RTC
The magazine of record for the embedded computing industry

Also in this issue

Network Processors

Shelf Management

Carrier Design for COM Express

February 2006 www.rtcmagazine.com

ATCA Offers Wide Integration and Implementation Options

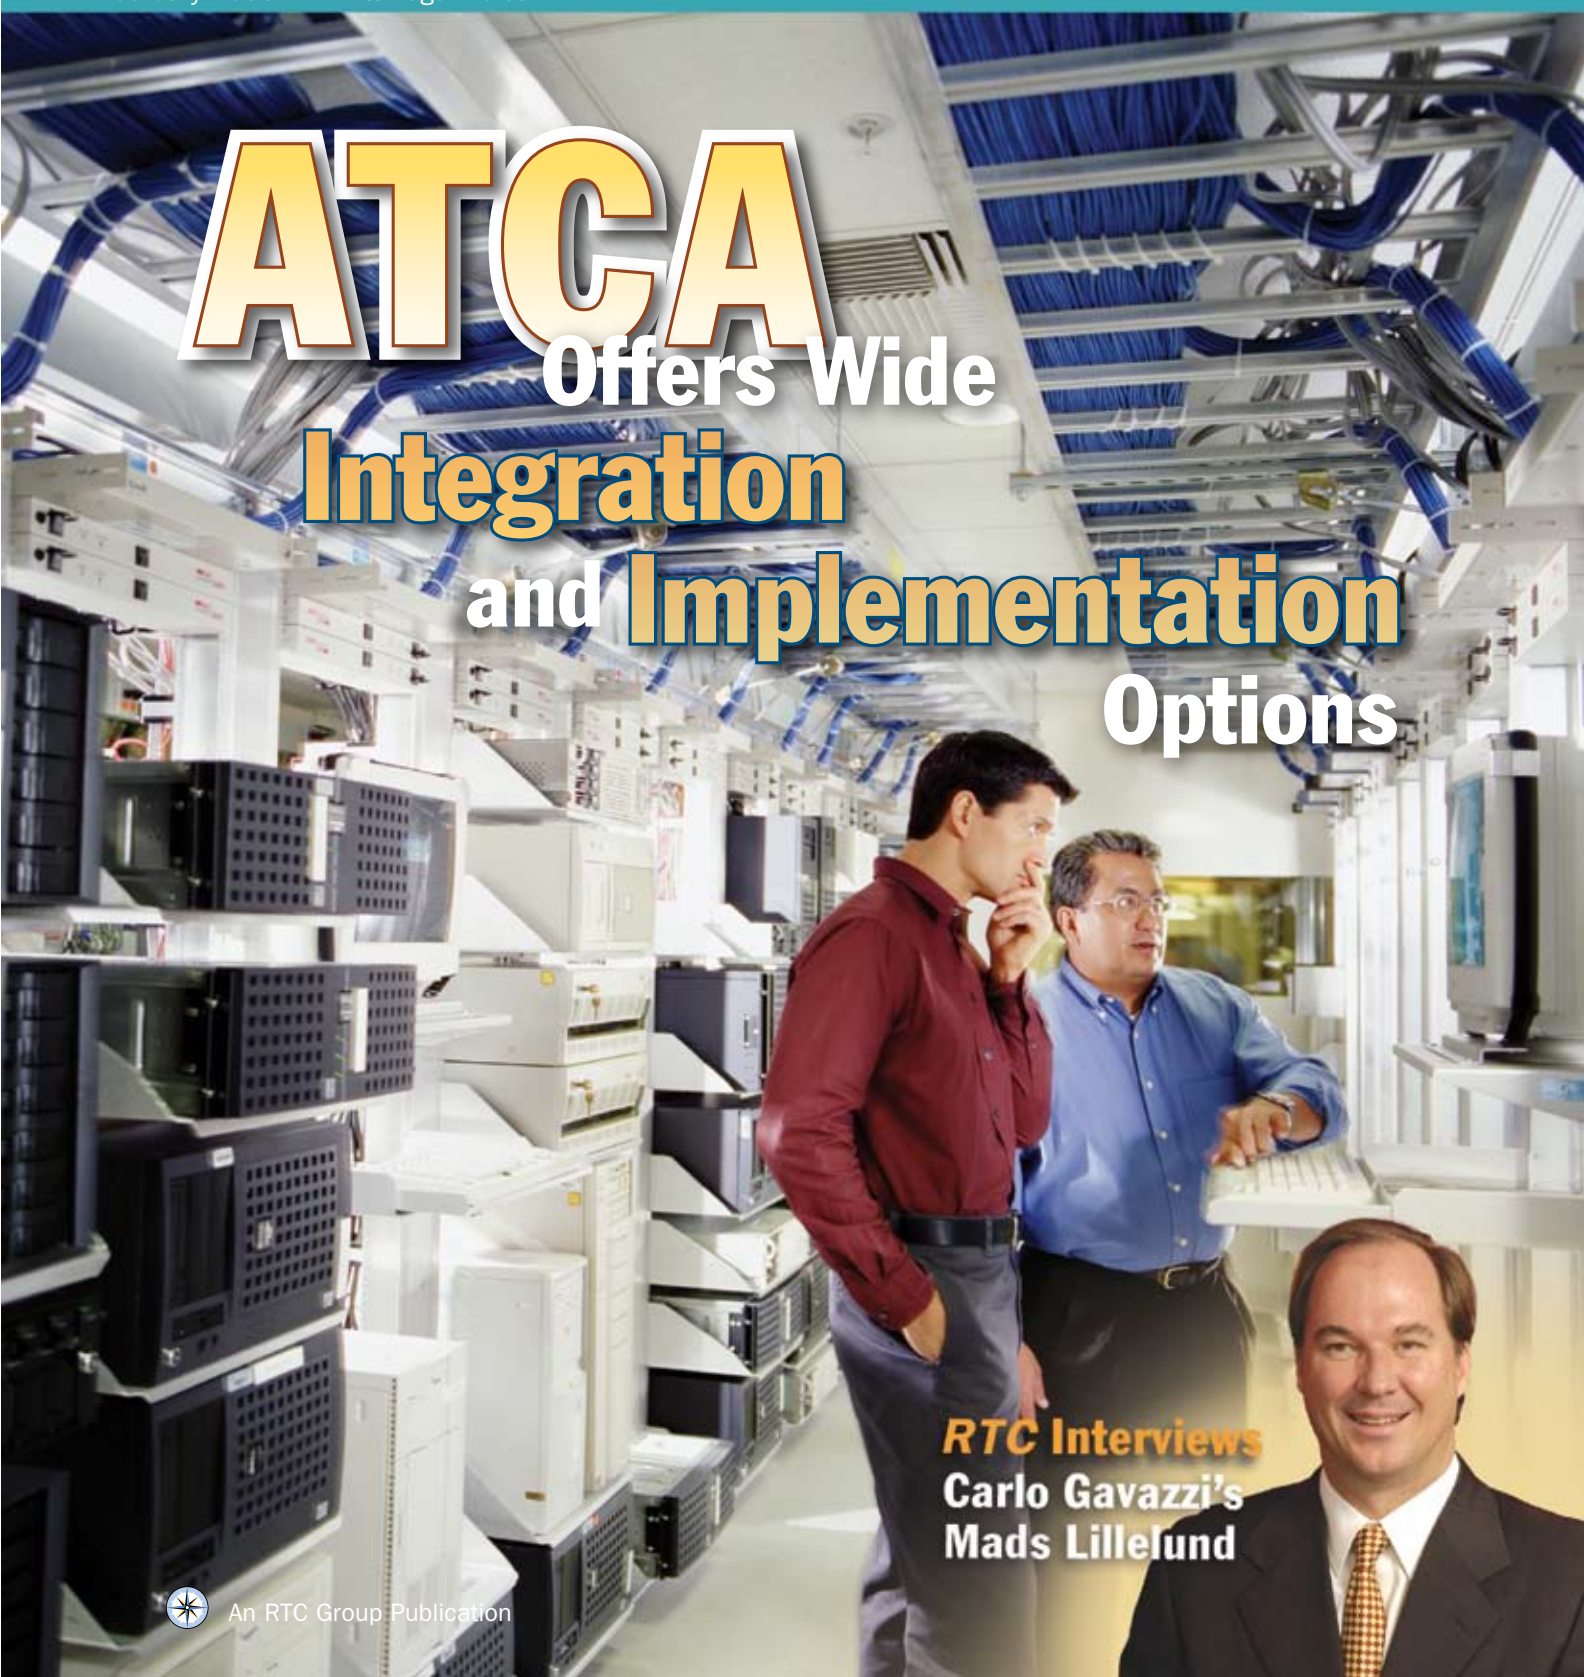

RTC Interviews Carlo Gavazzi's Mads Lillelund

An RTC Group Publication

Need reliability?
–40° to +85°

EPIC™ XE–900
1.0 GHz CPU

Features	XE–900	XE–800	XE–700
CPU	Via Eden	AMD Geode GXI	STPC
Clock speed	400 MHz; 733 MHz; 1.0 GHz	300 MHz	133 MHz
BIOS	General Software	Phoenix	Phoneix
DRAM support	to 256 MB	to 256 MB	32/64 MB
Compact/Flash	Type I or II	Type I or II	Type I or II
COM 1	RS–232	RS–232/422/485	RS–232
COM 2	RS–232	RS–232/422/485	RS–232/422/485
COM 3	RS–232	NA	RS–422/485
COM 4	RS–232	NA	RS–232
COM 5	RS–232/422/485	NA	NA
COM 6	RS–422/485/TTL	NA	NA
LPT1	0	0	1
EIDE	2	2	1
USB	2	6	2
CRT	1600 x 1200	1280 x 1024	1280 x 1024
Flat panel	LVDS	yes	yes
Digital I/O	24–bit prog.	48–bit prog.	24–bit prog.
Ethernet	10/100 Base–T	Dual 10/100 Base–T	10/100 Base–T
Expansion	PC/104 & *Plus*	PC/104 & *Plus*	PC/104
Power	3.6A operating	1.6A max.	1.6A max
Temp. range	–40° to 70/85° C	–40° to 80° C	–40° to 80/85° C
Shock/vibration	40/5g	40/5g	40/5g

Need Linux, QNX, Windows®?
Try our OS EMBEDDER™ KITS

Our kits are the shortest path to a successful OS on an Octagon embedded computer.

- Pick your Octagon SBC
- Pick the OS you prefer: Linux, Windows, QNX

Octagon delivers a high performance, total solution.

Typical Linux kit includes:

- *Target CPU card*
- *Preloaded OS image on 256 MB industrial CompactFlash*
- *256 MB SO–DIMM module*
- *Interface cables*
- *Hard copy of manual*
- *Mouse*
- *CPU OS bootable CD*
- *Optimized OS version*
- *Full driver support for on–board hardware*
- *X–Windows support*
- *Example applications and source code*
- *Extra documentation*

Need PC/104 expansion?
Try our XBLOKs®

XBLOKs offer the best compromise in cost and function for both PC/104 and PC/104-*Plus*. Only 44% the size of a standard PC/104 card, you can add two functions to your system but increase the stack height by only one level. –40° to 85° C. Heat diagram shows enhanced cooling.

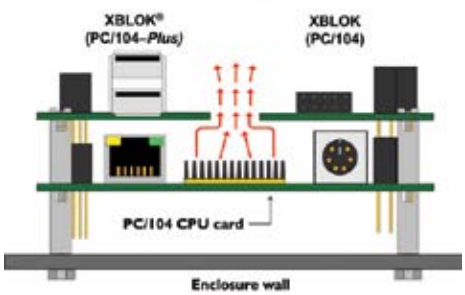

X–SRAM–2 MB
- 2 MB high speed, SRAM
- Read and write at full bus speed
- Pointers to memory saved if CPU resets or loses power

X–DIO–48 bit programmable digital I/O
- 48 digital I/O, 5V compatible
- Source and sink 16 mA per output
- Direct connection to opto-module racks

X–COM–2 dual UART
- Up to 230.4 kBaud data rate
- Supports RS–232/422/485
- RS–485 fault protected to ±60V

X–LAN–1 Ethernet LAN
- 10/100 Base–T, Intel 82551ER
- Fully plug–n–play
- High performance, PCI bus interface

X–USB–4 quad USB 2.0
- Speeds up to 480 mbps
- Mix and match USB 1.1 and 2.0
- Current–limited ports can supply 500 mA to external devices

Need a fanless system?
NEW CONDUCTION COOLING SYSTEM

Designed for the **XE-900**, our conduction cooling system eliminates a fan even at 1.0 GHz.

For a full listing of Octagon Systems products, visit us at
www.octagonsystems.com

OCTAGON SYSTEMS

That's an ATM to IP gateway?

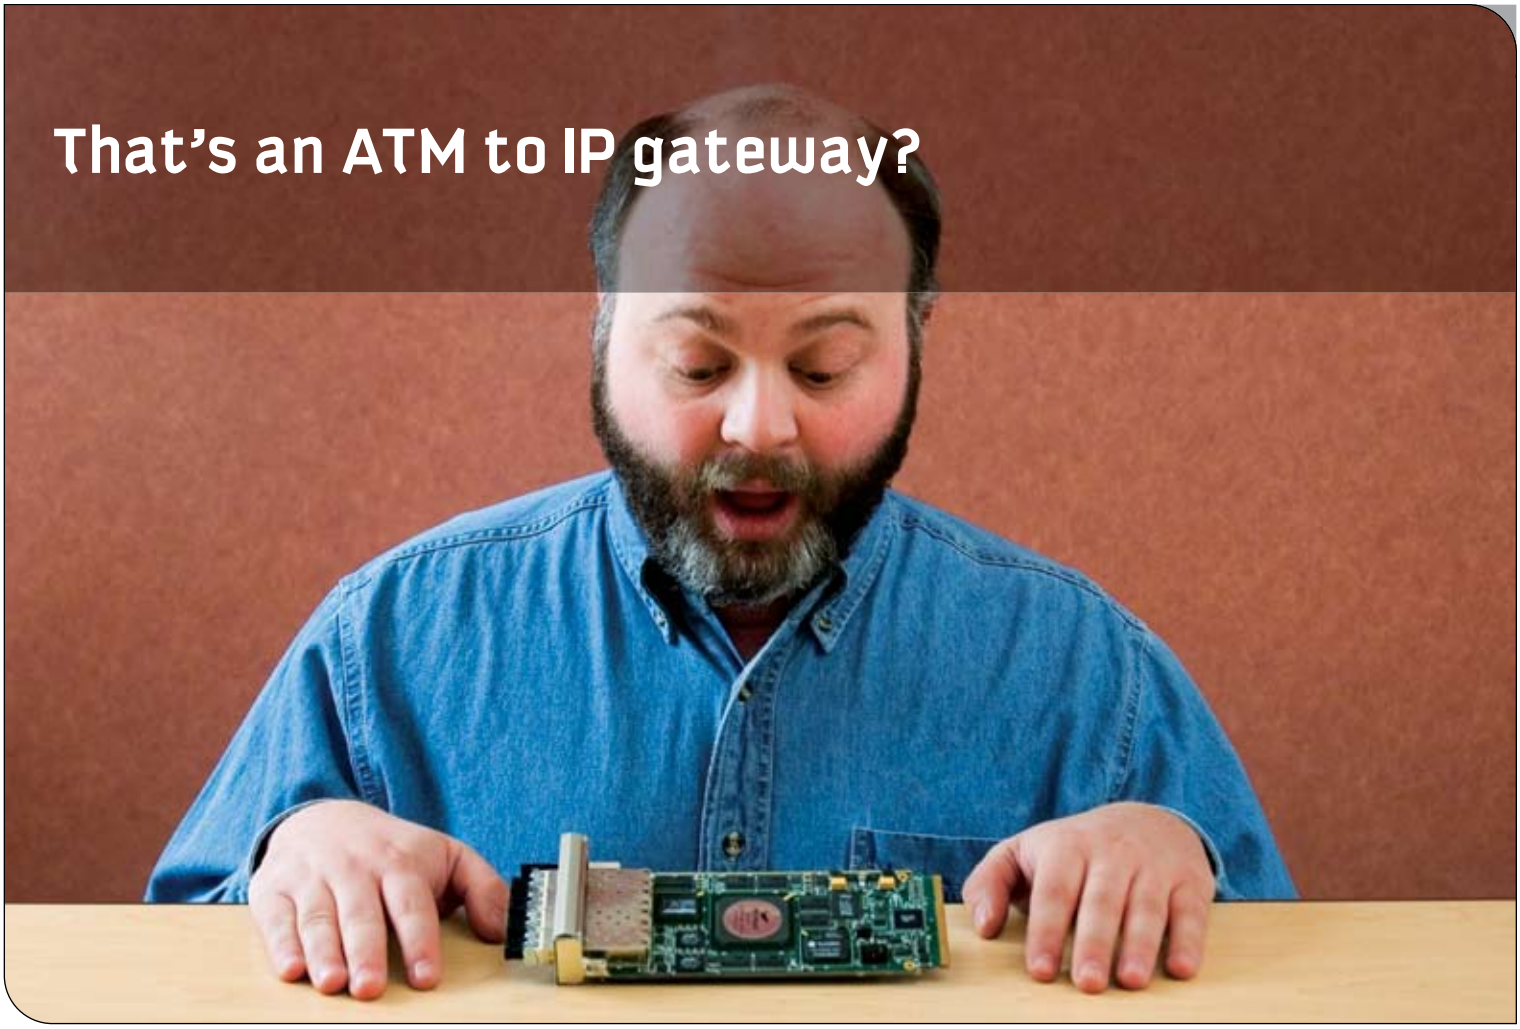

SBS knows AdvancedMC™. Get from ATM to IP on one board.

AdvancedMC™

Intel® Communications Alliance
Associate Member
SILVER

FITTING A GATEWAY that converts ATM to IP on an AdvancedMC™ wasn't exactly easy, but we did it. And while we were at it, we included the high availability features we knew you'd be looking for, like Automatic Protection Switching, Hot Swap, IPMI v1.5 for integrated management support and Carrier Grade Linux® driver support.

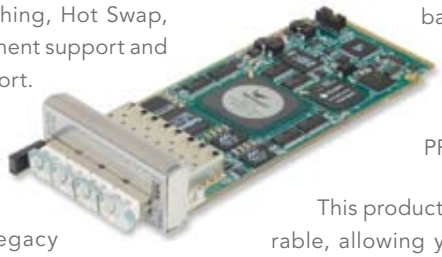

The Telum 1204-O3 is essentially a gateway in an AdvancedMC format specifically designed to facilitate the migration from legacy networks to IP-based networks. This board is ideal for any application where local Ethernet traffic needs to seamlessly access a wide area network using OC-3.

The Telum 1204-O3 module's versatile interworking capabilities allow it to independently interconnect between any supplied protocols, which can be selected on a per-port basis. For increased adaptability, a full set of IP services can be offered over any Layer 2 protocol including ATM AAL5, PPP and Ethernet.

This product is flexible and software configurable, allowing you to support a wide range of existing protocols as well as emerging standards. For a full description of features, please visit www.sbs.com/amc. Prepare to be amazed.

SBS Technologies

SBS knows.
Find the AdvancedMC board you're looking for at **www.sbs.com/amc** or call **800.SBS.EMBEDDED**

Departments

9 **Editorial:** Evolution vs. Intelligent Design

11 **Industry Insider**

56 **Products&Technology**

RTC

www.rtcmagazine.com

AudioCodes TP12610 ATCA media gateway board. • **Pg. 14**

Features

Technology in Context Enhanced ATCA

14 **Transitioning to VoIP: All-in-One Blades or Modular ATCA Solutions?**
Alan Percy, AudioCodes

20 **High Availability and Interoperability for Telco Class Applications**
Jeff Hudgins and Cynthia Tucker, Alliance Systems

25 **ATCA300: Expanding to New Applications**
Niall Robinson, Optovia

Solutions Engineering Network Processors

29 **Communications Processors vs. Network Processors: Programmable Data Plane Approaches**
David Smith, Freescale Semiconductor

40 **Beyond the Network Layer: NP-Based TCP Offloading via TCP Splicing**
Surekha Peri and Pravin Pathak, Agere Systems

Industry Insight Shelf Management

44 **Hardware Platform Management Building Blocks Speed ATCA and AMC Product Development**
Mark Overgaard, Pigeon Point Systems

48 **ATCA Shelf Management for High Availability**
Kevin Hart, Vadatech

Executive Interview

53 ***RTC* Interviews Mads Lillelund, CEO, Carlo Gavazzi**

Industry Watch

62 **Successful Carrier Card Design for COM Express**
Andrew Pitt, RadiSys

69 **Module-Based Design Platform Includes FPGAs**
Chris Wright, Ultimodule

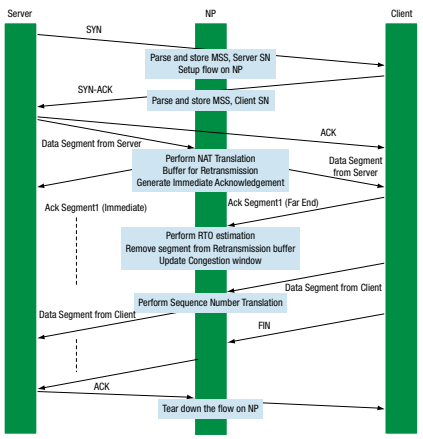

In TCP splicing, a well-known technique for enhancing TCP performance, two independent TCP connections span a session: one from the client to the NP, and the second from the NP to the server. TCP splicing on the NP-based TOE is shown here. • **Pg. 40**

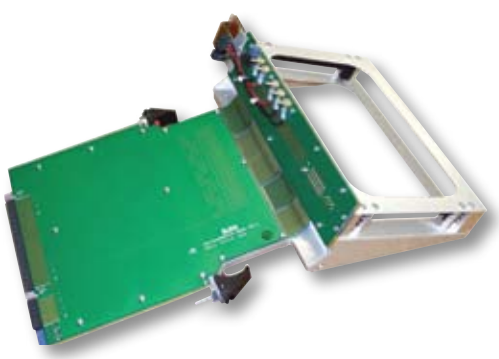

VXS Extender Board Provides Easy Access • **Pg. 56**

When Life Is On The Line
Operate and survive under the most extreme conditions with ruggedized E-Disk® solid-state flash drives and network storage solutions. BiTMICRO's cutting-edge storage technologies offer utmost reliability, optimum data security and unmatched performance.
Ethernet | Fibre Channel | SCSI | IDE/ATA
USB | FireWire | cPCI VME | SATA | iSCSI
PCI-X | PCI Express | SAS | Infiniband
BiTMICRO
ULTIMATE STORAGE SOLUTIONS
BiTMICRO Networks, Inc.
45550 Northport Loop E
Fremont, CA 94538-6481
www.bitmicro.com
info@bitmicro.com
510-743-3475

MISSION CRITICAL VME/cPCI
data storage modules
Extreme Comprehensiveness: We offer the most comprehensive VME/cPCI storage product line in the world, offering device alternatives for any standard or unique application.
• Solid State Disk • Removable Hard Disk
• Tape Drives • Optical Disk • PCMCIA Adapter
Extreme Performance: Our VME products feature extreme speed, capacity and ruggedly reliability with 320 MB/sec throughput enabled by LVD SCSI technology, storage capacity of more than 600 GBs per module and a 1,400,000 hour MTBF.
Extreme Quality: Phoenix International is the only manufacturer of VME data storage products that is ISO 9001:2000 Certified.
PHOENIX INTERNATIONAL
Phoenix International Systems, Inc. An ISO 9001:2000 Certified SDVOSB
714-283-4800 • 800-203-4800 • www.phenxint.com

Publisher

PRESIDENT
John Reardon, johnr@rtcgroup.com

VICE PRESIDENT, European Operations
Zoltan Hunor, zoltanh@rtcgroup.com

EDITORIAL DIRECTOR/ASSOCIATE PUBLISHER
Warren Andrews, warrena@rtcgroup.com

Editorial

EDITOR-IN-CHIEF
Tom Williams, tomw@rtcgroup.com

SENIOR EDITOR
Ann Thryft, annt@rtcgroup.com

MANAGING EDITOR
Marina Tringali, marinat@rtcgroup.com

STAFF EDITOR
Craig Choisser, craigc@rtcgroup.com

COPY EDITOR
Rochelle Cohn

Art/Production

CREATIVE DIRECTOR
Jenna Hazlett, jennah@rtcgroup.com

GRAPHIC DESIGNER
Melissa Gaeta, melissag@rtcgroup.com

PRODUCTION DESIGNER
Kirsten Wyatt, kirstenw@rtcgroup.com

DIRECTOR OF WEB DEVELOPMENT
Jan Garcia, jang@rtcgroup.com

WEB DEVELOPER
Marke Hallowell, markeh@rtcgroup.com

Advertising/Web Advertising

CALIFORNIA COASTAL ADVERTISING MANAGER
Diana Duke, dianad@rtcgroup.com
(949) 226-2011

WESTERN REGIONAL ADVERTISING MANAGER
Lea Ramirez, lear@rtcgroup.com
(949) 226-2026

EASTERN REGIONAL ADVERTISING MANAGER
Nancy Vanderslice, nancyv@rtcgroup.com
(978) 443-2402

EMEA SALES MANAGER
Anna Brown, annab@rtcgroup.com
(949) 226-2025

ASIAN ADVERTISING MANAGER
John Koon, johnk@rtcgroup.com
(858) 755-4999

Billing

Maggie McAuley, maggiem@rtcgroup.com
(949) 226-2024

To Contact *RTC* magazine:

HOME OFFICE
The RTC Group, 905 Calle Amanecer, Suite 250, San Clemente, CA 92673
Phone: (949) 226-2000 Fax: (949) 226-2050, www.rtcgroup.com

EASTERN SALES OFFICE
The RTC Group, 96 Dudley Road, Sudbury, MA 01776
Phone: (978) 443-2402 Fax: (978) 443-4844

Editorial Office
Warren Andrews, Editorial Director/Associate Publisher
39 Southport Cove, Bonita, FL 34134
Phone: (239) 992-4537 Fax: (239) 992-2396

Tom Williams, Editor-in-Chief
245-M Mt. Hermon Rd., PMB#F, Scotts Valley, CA 95066
Phone: (831) 335-1509 Fax: (408) 904-7214

Ann Thryft, Senior Editor
15520 Big Basin Way, Boulder Creek, CA 95006
Phone: (831) 338-8228

Published by The RTC Group

Copyright 2006, The RTC Group. Printed in the United States. All rights reserved. All related graphics are trademarks of The RTC Group. All other brand and product names are the property of their holders.

Open Architecture Solutions

Diversified Technology, Inc. (DTI) has a 34-year history of delivering embedded systems to the communications market and now builds on that by offering an increased line of performance-driven solutions based on the latest processing and switched fabric products. This line of AdvancedTCA (ATCA) systems, **DTI's Targa Series**, addresses the needs of the core and edge application markets, while also delivering unique features designed to handle the next generation services of advanced voice, data and mobile applications.

> Find out more at www.dtims.com/rtc

AdvancedTCA • CompactPCI • PCI/ISA • Embedded PCI-X • ETX • COM Express • Custom Design and Manufacturing

AdvancedTCA

Blade Platforms

- Complete AdvancedTCA Blade Systems packaged with Node Boards and Switch Blades (5 and 14 slot platforms - **DTI's Targa Series**)
- Multiple Node Boards and Switch Blades offered to provide increased functionality and customization features for Platform Solutions
- Fabric Agnostic Solutions based on Ethernet (PICMG 3.1) or InfiniBand (PICMG 3.2)
- Carrier Grade Linux
- Dual Star and Full Mesh Backplane Options

Embedded Systems Conference
San Jose, CA : **Booth 1928**

Diversified Technology®
An Ergon Co.

Three Decades of Embedded Solutions

CPU Boards - Switch Blades - Rackmount Systems - Modular Platforms
Computer Modules - System Integration - Custom Designs - Outsourcing Capabilities

1.800.443.2667 • sales@dtims.com • http://www.dtims.com/rtc

All trademarks and tradenames are the property of their respective owners.

The design possibilities are **wide open**

open

OPEN MODULAR SOLUTIONS

ACCESS

EDGE

CORE

TRANSPORT

DATA CENTER

AT8001 Processor

AT8902 Base/Fabric Hub

AM4500 AM4300 AM4001

XL8000

DESIGN AND DEPLOY

Your new IMS infrastructure applications using Kontron AdvancedTCA / AdvancedMC modular solutions.

Kontron simply takes the worry – and the expense – out of building complex IMS communication platforms for **fixed-mobile networks.** Whatever the application, you can design and deploy your project faster than you think with fully integrated, open standard modular solutions that are application-ready, right off the shelf.

That means reduced development costs for you, and tremendous **"swap in - swap out"** service flexibility for your carrier customers. It's a very smart win-win go-to-market strategy for everything from data and signaling platforms to IP streaming multimedia applications for video-on-demand, real-time voice and video telephony.

It's so simple. Start your next application with Kontron.

> Go Open Standards > Go Kontron > Ask for an Eval today >

www.kontron.com/ATCA

AdvancedTCA® AdvancedMC™ SERVICE AVAILABILITY FORUM ASI SIG Intel® Communications Alliance Premier Member

Computer On Modules | **Blades & Mezzanines** | CPU Boards | **Systems** | Mobile Rugged | Custom Solutions

kontron

1-888-526-ATCA - sales@us.kontron.com — EMEA: +49 0 800 7253756 - ASIA: +886 2 2910 3532

Kontron and the Kontron logo are registered trademarks of Kontron AG. All other trademarks are the property of their respective owners. ©2006 Kontron America, Inc.

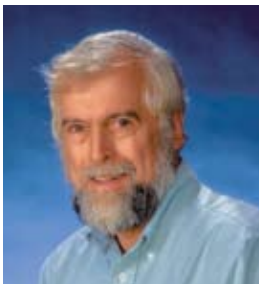

Evolution vs. Intelligent Design

by Tom Williams, Editor-in-Chief

Sometimes it doesn't hurt to draw back from the hurly-burly of day-to-day concerns and take a wider look at what is happening in our industry. By taking a look at how things have developed over the years, it may be possible to get a better understanding of how things will develop in the future. Sometimes it is astounding to see how far things have come and at other times it is astounding that anything works at all.

To borrow a controversy from biology: Is this all the work of an intelligent designer, or did (and does) it all evolve at random? In our case, I think the answer is both. In the world of digital electronics there are literally many thousands of intelligent designers. In the world of technology, competition and markets, there is definitely a process of evolution and some of it even seems random.

Is this all the work of an intelligent designer, or did (and does) it all evolve at random?

At the recent Bus & Board conference in Long Beach, we were treated to a slide show giving a 25-year retrospective of VITA. Younger versions of familiar faces flashed across the screen, along with long-lost company names and pictures of products and engineers at test equipment designing things the world has long forgotten, but which were considered vital at the time.

I was struck by the thought that for every product put on the market, no matter how small, there have been countless meetings between management, marketing and engineering groups. There has been design and test, prototyping, manufacturing and more testing. There have been marketing strategy meetings, promotional and sales materials, training classes, sales meetings, (hopefully) advertising, shipping and (hopefully) customer support. And that's just for the products that actually make it to market.

As an editor, I've been told (and written) about a vast number of products and ideas that never made it that far. Some of the products that made it were successful and became the beginning of whole product lines, families and sometimes standards. For these I got credit for insightful reporting. Others simply died a quiet death and either caused companies to regroup or go under. For these I got called a chump. Interestingly, most of them seemed like good and intriguing ideas at the time. Is there a pattern to all this?

There must be some kind of semi-Darwinian mechanism at work here because technology has made amazing advances in the past 20 years. Certainly, in the big picture it is difficult to discern the guiding hand of an all-knowing intelligence. If that's what market research is looking for, it is probably doomed to failure. In the last analysis, it appears that legions of highly skilled and intelligent researchers and designers are grasping about in semi-darkness and making incredible progress partially in spite of themselves, bumping into things and despite it all struggling forward.

Predicting the advance of technology is a little like predicting the weather. Using vast amounts of data, it is possible to make fairly reliable predictions about five days out. Beyond that, reliable analysis breaks down because it is essentially a chaotic system. And so we continue on, and that's what makes observing this industry—and participating in it—such an adventure.

Leading the way in Digital Receiver Technology

When only the best is good enough.

Delivering the best solution means not making compromises. Not compromising performance – when you could have the industry-leading sustained performance of the ICS-8552 or ICS-8554. Not compromising reliability – when you could have the world-beating expertise of ICS in developing rugged solutions. Not compromising choice – when the ICS-8552 and ICS-8554 provide the best in ADC, DDC and FPGA technology for software defined radio, military communications, radar and signal intelligence applications. And not compromising your budget when you can choose the optimum solution for either development or deployment.

Or is second best good enough?

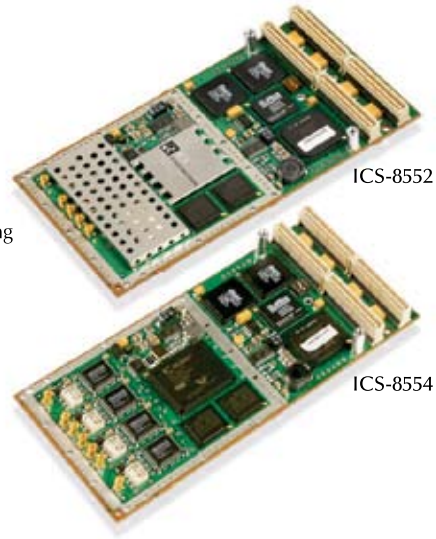

ICS-8552

ICS-8554

ICS
SENSOR PROCESSING

CHECK THE FULL SPEC NOW AT
www.ics-ltd.com/8554
www.ics-ltd.com/8552

OR CALL TOLL FREE (US ONLY)
800-267-9794 or 613-749-9241

INDUSTRY INSIDER

VMetro Scores Gains with BAE Systems and NASA

VMetro has received a contract to supply digital signal processing boards to BAE Systems in Nashua, NH for the Airborne Reconnaissance Low (ARL) Communications Intelligence (COMINT) Upgrade program, a two-year effort to provide an upgraded COMINT capability to the U.S. Army ARL fleet. VMetro's Phoenix VPF1 Dual PowerPC, Dual Virtex-II Pro FPGA VME/VXS Digital Signal Processors and PMC-FPGA02 Xilinx Virtex-II FPGA PMC Modules will be contained within the system provided to the Army for installation aboard the ARL platform. The VPF1 and PMC-FPGA02 boards will run the signal processing algorithm as well as perform real-time processing of the sensor data. This initial contract is valued at approximately $850K. Along with this award, this design is being included in similar future programs as part of BAE Systems' common platform initiative.

Additionally, NASA has presented VMetro with a Certificate of Recognition for being a quality supplier to NASA and for supporting many NASA robotic programs. During the visit, NASA Supplier Outreach and Process Control Assurance (SOPCA) representative, Chuck Holloway, met with VMetro management and employees to review several NASA initiatives and gave a presentation on Process Control Awareness. SOPCA is a NASA-wide effort to provide suppliers with information and tools to aid in the continual improvement of process control.

Xilinx Acquires DSP Design Tool Company AccelChip

Xilinx has announced that it has acquired AccelChip, a provider of Matlab synthesis software tools for building digital signal processing (DSP) systems. The AccelChip DSP Synthesis tools and DSP library of algorithmic intellectual property (IP) will be offered as part of Xilinx XtremeDSP solutions that include the System Generator for DSP design tools, DSP libraries and reference designs. This combination of DSP technologies enables thousands of DSP algorithm and system designers who use Matlab and Simulink from The MathWorks, to create systems using Xilinx reconfigurable DSPs.

Consistent with the DSP strategy and product roadmaps unveiled by Xilinx in October 2005, tools such as those offered by Xilinx and AccelChip accelerate the adoption of application-optimized solutions for digital communications; multimedia, video and imaging (MVI); and defense systems. Together the companies have significant customer overlap and synergy across these market segments, key elements for targeting the $2 billion high-performance DSP market.

Financial terms of the acquisition were not disclosed, however, the acquisition represents less than one percent of the total assets of Xilinx as of the end of its third quarter ended December 31, 2005. The company anticipates booking the acquisition in the March 2006 ending quarter. The majority of AccelChip employees, located at the Company's headquarters in Milpitas, CA and two design centers in North America, will become employees within the Xilinx DSP Division.

Xilinx's strategy is to deliver application-optimized XtremeDSP solutions on the industry's most compelling performance per dollar reconfigurable DSP using Xilinx programmable devices. The XtremeDSP solutions encompass specialized hardware development platforms, algorithmic libraries, as well as robust design and verification tool suites.

Systems Modeling Language Standard Submitted to the OMG

The SysML Submission Team (SST), chaired by Sanford Friedenthal of Lockheed Martin Corp., has submitted its proposed SysML specification to the Object Management Group (OMG) for review. The SysML v0.98 specification submission is in response to the joint Request for Proposal issued by the OMG and the International Council on Systems Engineering (INCOSE) for a customized version of UML 2.0 designed to address the specific needs of system engineers. With the submission of the specification, the SysML language will enter the OMG's final standardization process.

SysML is a new visual modeling language that extends UML 2.0 in order to support the specification, analysis, design, verification and validation of complex systems that include components for hardware, software, data, personnel, procedures and facilities. SysML reuses a subset of UML 2.0 diagrams and augments them with some new diagrams and constructs appropriate for systems modeling. SysML is designed to complement UML 2.0, so systems engineers who are specifying a system with SysML can collaborate efficiently with software engineers who are defining a system with UML 2.0.

The completion of the SST's SysML v0.98 specification is the result of cooperation between some of the foremost tool vendors, leading industry users, government agencies and professional organizations. Vendors, including ARTiSAN Software Tools, I-Logix, EmbeddedPlus Engineering, IBM, Mentor Graphics and Sparx Systems, closely collaborated to create the specification. Other stakeholders, including American Systems Corporation, BAE SYSTEMS, The Boeing Company, Thales, EADS Astrium, Eurostep Group, Lockheed Martin Corporation, the National Institute of Standards and Technology, oose.de Dienstleistungen für innovative Informatik, and Raytheon Company, added their own expertise to the specification effort, making it the most comprehensive, collaborative standards definition effort to date.

Carlo Gavazzi and Comtel Team for ATCA Development and Marketing

Carlo Gavazzi Computing Solutions and Comtel Electronics GmbH have joined into a strategic technology partnership to develop and market open architecture chassis, backplanes and shelf managers. The partnership leverages both companies' engineering resources and technologies to address the emerging ATCA and other open architecture markets in both North America and Europe.

The first collaborative effort generated by the partnership yields a line of AdvancedTCA (ATCA) products including 2U two-slot through 13U 16-slot shelves. Each system may be configured with full mesh or dual star backplanes and Pigeon Point or Motorola shelf management solutions. The modular and flexible set of offerings can be configured and integrated to support development and proof of concept systems through production-ready RoHS and NEBS deployable platforms. The partnership also allows for efficient scalability for medium

Get Connected with companies mentioned in this article.
www.rtcmagazine.com/getconnected

to high volume applications. Other open architecture products are being explored for mutual development and marketing between the two companies.

According to Mads Lillelund, president and CEO of Carlo Gavazzi Computing Solutions, "Our combined sales, engineering and manufacturing capabilities allow us to quickly design and develop best-in-class ATCA products as well as localize design services for our customers." Berthold Michels, president of Comtel Electronics GmbH added, "This technical collaboration allows us to leverage Carlo Gavazzi's continuous participation in the open standards community as well as broaden our product offerings within the ATCA ecosystem."

Kontron Converts Order System to Lead-Free/ RoHS

From January 2006, Kontron has announced its order system for board-level products will be converted to become RoHS-compliant on an order-related basis. Lead-free products complying with RoHS will thereby become standard, and solutions containing lead will only be available as special customer-specific orders. At the same time, existing customers for whom RoHS-compliant products cannot be delivered will be informed about redesign options. This mainly involves older solutions with the ISA bus that are affected. Thanks to Kontron's PCI-to-ISA bridge that can be soldered without using lead, redesigns for ISA will still be possible. For VME as well, lead-free solutions of every commonly used performance class are available. Detailed RoHS product information, including the lead content of all products, was supplied in Kontron's July 2005 product catalog. On Kontron's Web site, customers will find comprehensive information regarding the issue of lead-free/ RoHS (www.kontron-emea.com/ leadfree). A special e-mail contact form is also provided where they can inquire about the status of the products they use, including those that are customer-specific.

All system solutions will also be RoHS-compliant by June 30th, 2006. Some details still need to be worked out however, for example, in front coatings as well as in related display technology. Immediately after the lead-free/RoHS directives were adopted in 2003, Kontron introduced an RoHS working group to ensure that all Kontron board-level products would be lead-free by January 2006. According to the company, however, conversion to lead-free production is not simply a legal requirement, but also an active contribution to a lasting, environmentally friendly business.

NextHop and AirDefense Partner to Deliver Wireless Security

NextHop Technologies has announced a partnership with AirDefense to jointly provide their solutions that enable enterprise OEMs to integrate AirDefense's rigorous wireless security for mobile employees with NextHop UNS 4.0 unified switching software for differentiation and faster time-to-market.

NextHop's Dynamic Network Architecture (DNA) delivers modular software components that are portable across industry standard silicon and real-time operating systems (RTOS). AirDefense uses the DNA APIs for easy integration of its value-added application with unified management and control interfaces for an easy to understand and operate system. AirDefense provides uninterrupted protection for all mobile employees and their enterprise wireless assets, regardless of location—at work, home, airports or other wireless hotspots.

AirDefense Personal can easily be installed on a user's laptop and within minutes allows the user to monitor the laptop for potential wireless security risks. When integrated with AirDefense's flagship product, AirDefense Enterprise, the combined solution enables network administrators to effectively protect information assets automatically from threats not only within the four walls of the enterprise but also outside the confines of the building. Policy profiles are defined centrally on AirDefense Enterprise and automatically downloaded to each mobile user. If threats are discovered, AirDefense Personal notifies the user and sends the alerts to AirDefense Enterprise for central reporting and notification. This unique solution allows the network administrator to enforce corporate policies and provide complete protection for the mobile worker, regardless of location.

Radstone Forms Systems Applications Integration Business

Radstone Embedded Computing has announced the formation of a new Systems Applications Integration business unit. The new unit is tasked with developing customer solutions that leverage the breadth and depth of the technology, expertise and services that exists within the organization.

The Systems Applications Integration business unit will feature a network of Application Centers to enable "front line" development and support activities across the North American and ROW territories. The Applications Center located at Radstone's North American headquarters in Billerica, MA, will comprise a cross-functional team of application-focused engineers and support staff, and will be the primary point of contact for all U.S. customers requiring Radstone solutions that cannot readily be provided from the standard product catalog—notably, solutions in mission computing, radar and sonar.

Radstone has previously announced its strategic commitment to developing both VME and VPX (VITA 46) products, to integrated bridge processor architectures such as the MPC 8641 and to the 3U form-factor. The company expects to announce, nearer the date of planned shipment, a range of VPX single board computers—including multiprocessor products—graphics processors, Ethernet switches and Software Defined Radio blades based on the 3U form-factor. Radstone expects that the 3U form-factor will become increasingly important in the future as a result of the rapidly increasing demand for systems that are smaller, lighter and consume less power—and that integrated bridge processor architectures will deliver the functionality density to allow 3U solutions to achieve levels of performance previously only thought possible with 6U solutions.

Event Calendar

02/28/06
Real-Time & Embedded Computing Conference
Melbourne, FL
www.rtecc.com/melbourn

03/02/06
Real-Time & Embedded Computing Conference
Huntsville, AL
www.rtecc.com/huntsville

03/05-10/06
Optical Fiber Comm. Conference & Expo
Anaheim, CA
www.ofcnfoec.org

03/07-09/06
Intel Developer Forum
San Francisco, CA
www.intel.com/idf

03/14-15/06
Military Technologies Conf.
Boston, MA
www.miltechconference.com

03/20-23/06
Nat'l Manufacturing Week
Chicago, IL
www.manufacturingweek.com

If your company produces any type of industry event, you can get your event listed by contacting sallyb@rtcgroup.com.
This is a FREE industry-wide listing.

"The QNX Neutrino microkernel OS has the perfect DNA for multi-core processors"

Dan Dodge. QNX CEO & CTO.
Pioneer in distributed and multiprocessor computing.

Introducing the **QNX® Momentics® development suite Multi-Core Edition,** the industry's most comprehensive software platform for multi-core systems. Powered by the massively scalable QNX Neutrino® RTOS, this fully integrated solution supports AMP, SMP, and BMP, a groundbreaking technology that simplifies code migration and future-proofs your designs for quad-core and beyond. It's the latest innovation from QNX Software Systems, the undisputed leader in multiprocessing technology.

Maximum Choice for Multi-Core

Only QNX gives you the power to choose the best multiprocessing model for your multi-core design:

	SMP	BMP	AMP
Seamless resource sharing	✓	✓	—
Scalable beyond dual-core	✓	✓	Limited
Mixed OS environment	—	—	✓
Dedicated processor by function	—	✓	✓
Inter-core messaging	Fast (OS primitives)	Fast (OS primitives)	Slower (application)
Thread synchronization between cores	✓	✓	—
Dynamic load balancing	✓	✓	—
System-wide debug & optimization	✓	✓	—

QNX Unlocks the Power of Multi-Core

Maximize performance. Eliminate complexity. Accelerate migration. Only QNX offers:

- **Asymmetric Multiprocessing (AMP)** for full developer control and fault tolerance
- **Symmetric Multiprocessing (SMP)** for maximum concurrency and scalability
- **Bound Multiprocessing (BMP)** for the fastest code migration and minimum design complexity
- **Transparent Inter-Processor Communication (TIPC)** protocol for seamless Linux connectivity
- **System tracing tools** for fast debugging and optimization of multi-core applications
- **Off-the-shelf BSPs** for multi-core platforms based on MIPS®, PowerPC®, and x86 architectures

Discover how Dan and the QNX team deliver the shortest migration path to multi-core. Call 1 800 676 0566 or visit **www.qnx.com/innovate.**

QNX SOFTWARE SYSTEMS

QNX, Momentics, and Neutrino are trademarks or registered trademarks of QNX Software Systems GmbH & Co. KG and are used under license by QNX Software Systems International Corporation. All other trademarks belong to their respective owners. 301786 MC339.05

Transitioning to VoIP: All-in-One Blades or Modular ATCA Solutions?

The combination of the ATCA form-factor along with the debut of the AMC modular card offers developers a wide choice of options for implementing network and communication equipment—to integrate many functions on one board or to customize a larger carrier board with a selection of modules?

by Alan Percy
AudioCodes

One hundred years of technology investment has created one of the most reliable voice networks the world has seen. But the writing is on the wall: the legacy infrastructure is soon to be replaced by more agile Voice-over-IP technologies, carrying voice as packets over a robust real-time data network. From the tried and true ILECs to the latest "download me now to make free phone call" challengers, the race is on to upgrade or replace legacy equipment.

This creates an enormous opportunity for developers who understand the needs, priorities and solutions available to create Voice-over-IP equipment to replace the legacy infrastructure. The developer's challenge is figuring out which of the competing form-factors provides the right platform to serve as a foundation for the next-generation network.

Shifts in Technology

A number of changes in technology have fueled a shift in the way network equipment is developed. The evolution of the newly defined AdvancedTCA form-factor and its offspring AMC modules along with MicroTCA form-factors has provided a framework that enables multiple vendors' blades to interoperate in a highly reliable and mechanically sound package. Along with these standard form factors a set of standard protocols has arisen. Now, instead of closed protocols required for command and control of network components, standard protocols, including SIP, H.323 and MGCP, allow different vendors to interoperate their equipment without expensive custom development.

On all fronts, advances in processor technology have brought intelligence down to the module level. Increased computing capabilities in smaller packages and improved embedded operating system technologies have created a shift to AMC module-level intelligent devices. Complete media gateway

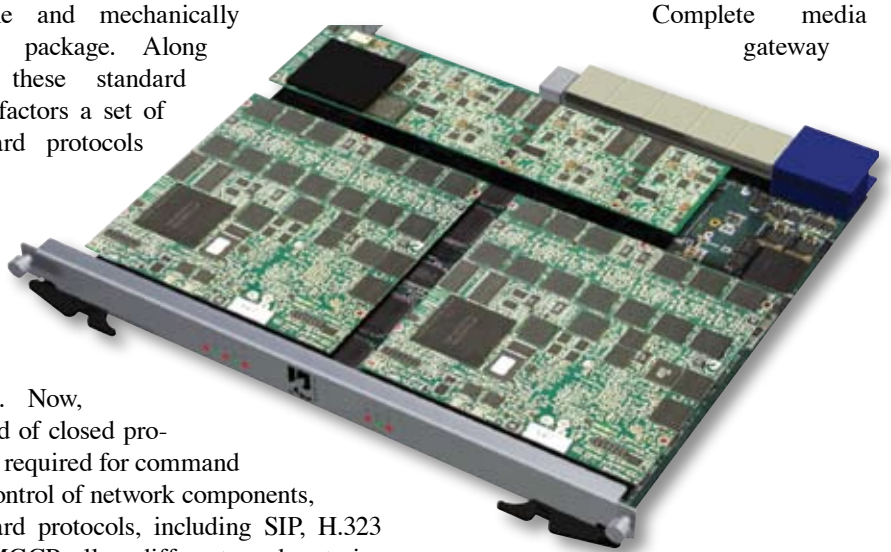

Figure 1 AudioCodes TP12610 ATCA media gateway board.

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

and media server technologies are now available as modular one-blade resource cards. With these changes in technology, a whole new paradigm has evolved that allows integrators to build not just at the blade level but also at the module level.

The ATCA specification was designed to help manufacturers lower equipment cost by permitting the use of off-the-shelf components. The ATCA also increases flexibility and functionality by providing a large board area—8U by 280 mm deep—and a 200-watt power draw (Figure 1). This accommodates dense, processor-intense designs that meet the demand for new high-speed data services and expanded carrier IT services. The combination of functionality, flexibility and redundancy of the ATCA specification produces a platform that carriers can use to optimize services and improve profitability.

Leveraging Blade and Module-Level Resources

Designing telecommunications solutions at the blade level typically includes choosing one or more standard form-factor ATCA boards (or blades) that perform some portion of the total solution functionality. These blades are combined into an off-the-shelf chassis to create an integrated application or network element. Examples of complete blade-level components are:

- CPU Blades – containing complete Pentium-class processors, memory and hard disk on a single board able to host complex applications.
- Media Gateways – providing both network connectivity and media conversion needed to interface between the legacy TDM network and VoIP-based networks
- Media Servers – offering complex media processing capabilities that include transcoding, conferencing, tone detection, etc.
- PSTN Interface cards – providing network interfaces and signaling to the legacy TDM network using CAS, ISDN or SS7 signaling

Building solutions at the blade level is not new. In fact, designers have done similar integrations for years. What is new are the standards that make integration easier and less costly. Standard-based protocols and the ATCA base interfaces have replaced custom- developed drivers and APIs that were hard to use and even harder to debug. Intricate dependencies between the host operating systems, device drivers and APIs have all but been eliminated in standards-based solutions. As a result, the ability to build custom solutions at the blade level is now within reach of many more engineers.

One of the unique features of the ATCA form-factor is formalization of the ATCA Mezzanine Card (AMC), a module-level component that can be added or combined with other modules into one ATCA slot (Figure 2). In many cases, a full-size ATCA blade with native base functionality (e.g., a CPU blade) can provide one or more open AMC module slots. In this situation, an AMC slot permits flexibility in adding expansion or additional functionalities into the existing blade. Adding hard disk storage is one example of an ideal application for AMC modules on a CPU blade (Figure 3).

In other cases, a designer may find all or many of the needed resources on one or more AMC modules and would like to integrate these resources together into one ATCA slot. When this is the case, an AMC carrier blade would allow the various modules to be combined into a custom ATCA blade slot. AMC carrier blades have no native functionality; they simply provide connectivity, power and cooling to the associated AMC modules.

Making a Decision

With so many choices between full-size ATCA blades, leveraging AMC module add-ons or complete AMC-based solutions, what are the advantages of using one approach vs. another?

Full ATCA Blade-level integrations offer a number of advantages that designers find appealing. Full-size ATCA blades can access the optional rear transition module (RTM). In highly specialized network interface technologies, such as T1/E1 interfaces, VoIP gateways and ATM, the design of the blade is generally split to leverage the RTM feature found on ATCA blades. This brings the network interface logic to the rear of the chassis, simplifying wiring and serviceability. By focusing the technology design into one dedicated blade, vendors are also able to optimize the density of the product, utilizing every square inch of the ATCA form-factor.

While ATCA remains a mechanically complex platform, building a one-blade solution reduces the number of connectors and mechanical support needed for a number of AMC modules with positive effects on reliability. ATCA supports sophisticated redundancy and High-Availability (HA) features that allow one

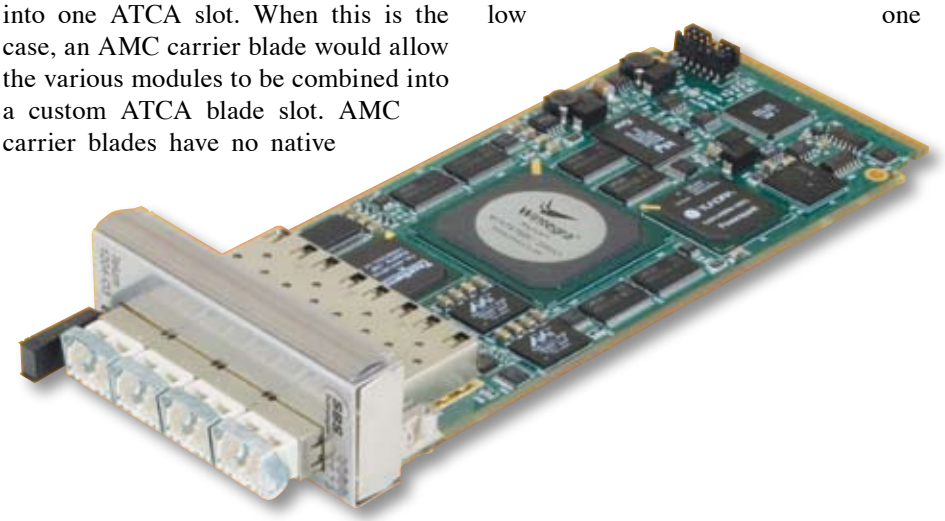

Figure 2 A network processor-based full-height Telum 1204-03 AMC module from SBS designed for high-availability telecom edge access applications.

Situation	Suggested Path
Project requires interfaces to existing PSTN or other legacy physical interfaces. Based on requirements, legacy interface and High Availability are top priorities.	In most cases, legacy PSTN interfaces will require a Rear Transition Module (RTM). Only full-size ATCA blades support RTM interfaces and can provide the rear connectivity required in carrier-class applications.
Project has legacy PSTN interface requirements but no single vendor can deliver a complete one-slot solution. High Availability is not needed.	Investigate the potential of an ATCA blade with open AMC slots that can be populated with a set of modules to complete functionality.
Project has no legacy PSTN interface requirements but does require a range of processing and storage needs.	Consider an ATCA processor blade with optional AMC slots to allow expansion that might include DSP processors, disk storage or other complementary technologies.

Table 1

THE COMPLETE DEVELOPMENT SOLUTION
Now offering Linux OS aware debug for ARM-based processors! Call us today!
Arium offers robust JTAG emulation and development tools for today's embedded software engineers using targets with ARM7™/ ARM9™/ARM11™, Intel XScale®, and TI OMAP™ cores and Intel® Pentium® processor families.
• Full symbolic, source-level Linux kernel debug and source-level process debug; seamless debug between them! No other vendor offers this powerful feature at any price.
• Real-time, integrated ETM trace data collection at 640 MHz and a GByte of trace memory.
• Highly integrated SourcePoint™ IDE with powerful, flexible code editing with debug integration.
• Real-time performance analysis for faster, more accurate results.
• Fast, easy, intuitive run control with robust C-like command language facilities.
• SourcePoint debugger available for Microsoft® Windows® and Linux hosts.
It's hard to compete without the right tools!
american arium
American Arium • 14811 Myford Road • Tustin, CA 92780 • 877.508.3970 • www.arium.com

blade to take over when another one fails. The complexity of HA often prevents multiple vendors from integrating these features at the AMC module level. In addition, if a single vendor designs the blade, the probability of inter-vendor compatibility issues is reduced.

However, many solutions can leverage a combination of base functionally on an ATCA blade and one or more AMC modules. This approach also offers some unique advantages, one of which is flexibility. Developers can now choose from a wide array of special-purpose AMC modules, enabling new combinations of features in the same shelf. This often lets designers eliminate separate blades or whole chassis from the system design, saving space and money.

And let's face it. No single vendor can do it all. Integrating at the AMC level also opens the opportunity to use multiple AMC vendors in one slot to get the unique mix of performance and functionality that can add value to a design. By stocking

Figure 3 Kontron ATCA CPU blade with optional AMC modules.

module-level AMC components, failures in the field can be quickly repaired or replaced without removing or replacing an entire ATCA blade. Also, stocking only AMC modules reduces the cost of maintaining a spares inventory.

Making the Choice

Given the array of advantages of both blade-level integrations and pre-packaged solutions, how does a developer choose a path forward? Table 1 offers some possible situations in which a company might find itself and offers a suggested development path.

As the legacy carriers and next-generation communication tools begin to leverage VoIP technologies for infrastructure, ATCA and AMC modules will surely play an important role. By choosing the right combination of blades and module technologies, developers have an opportunity to maximize the value proposition for their customers and create a network that will continue to provide reliable services for the next 100 years. ◢

Looking For More?
Visit www.rtcmagazine.com to download additional technical information related to this article.

AudioCodes
Somerset, NJ.
(716) 662-2769.
[www.audiocodes.com].

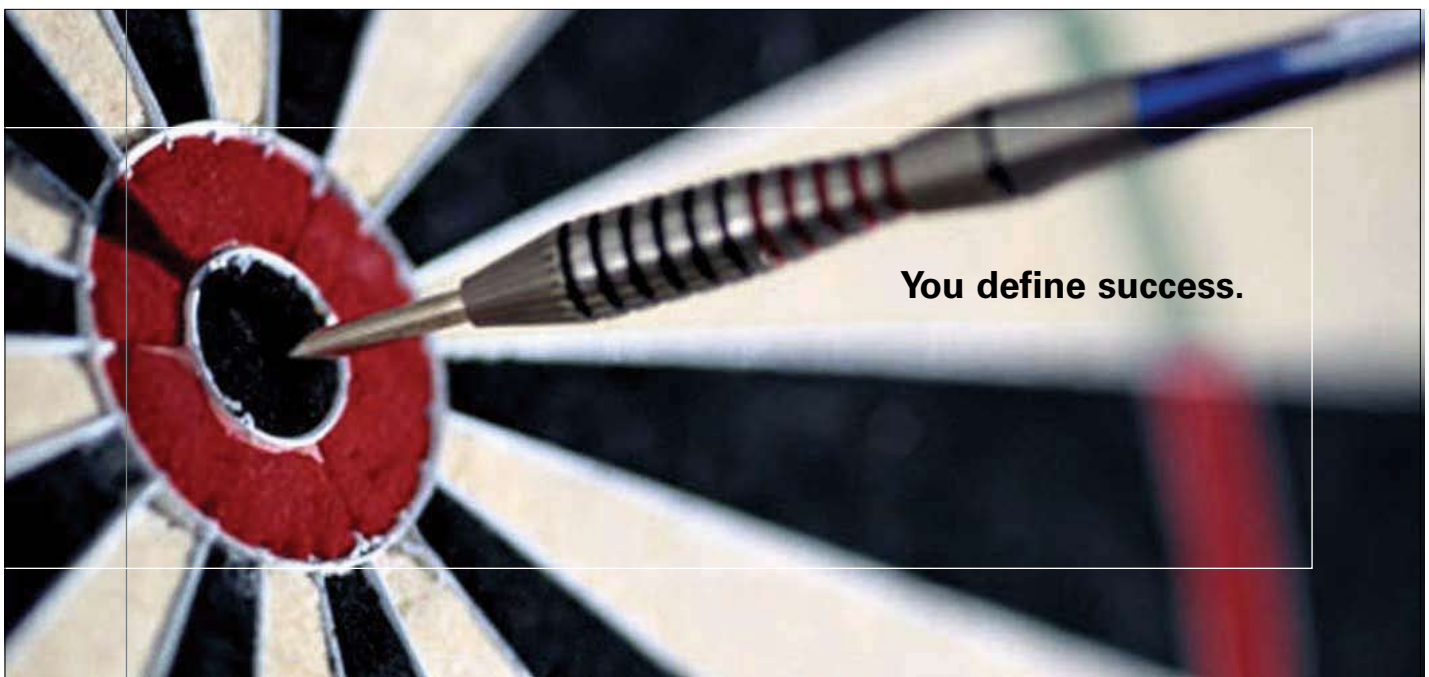

We design solutions.

Mercury customers span multiple industries and face unique computing challenges – whether improving yields for semiconductor wafer inspection, increasing throughput in medical imaging, rendering high-quality animation within a defined budget, or packing enormous processing capacity in deployed ground vehicles. Why are we so driven to tackle these difficult computing problems? Because at Mercury your challenges drive our innovation.

Our new PowerBlock™ 200 system incorporating Cell technology delivers unprecedented processing density within a rugged ATR chassis for vetronics applications. Contact us to learn how our products and services can optimize your challenging applications.

Let us design an innovative solution for you.
www.mc.com/rtc2 | 866-627-6951

Computer Systems, Inc.
MERCURY
Challenges Drive Innovation™

Perform like never before.

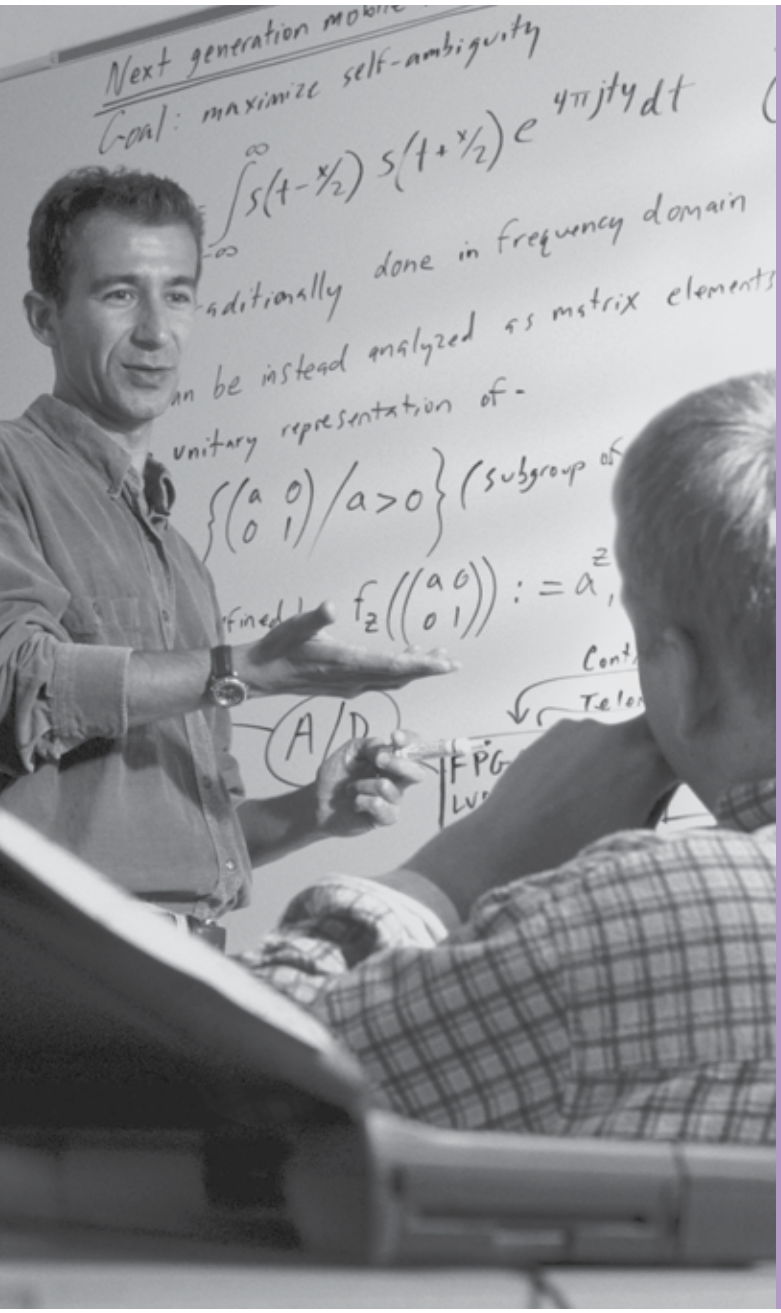

ANALOG DEVICES SHARC®

ANALOG DEVICES TigerSHARC®

From $5 to 3.6 GFLOPS, your DSP job just got easier.

Worried about rounding errors in your 16-bit arithmetic? Anxious about noise in your frequency domain processing? What about blowing your project schedule as you move your floating-point concept to fixed-point implementation? Analog Devices SHARC® and TigerSHARC® Processors solve your most frustrating numeric DSP challenges. Now you can really perform.

For more than 10 years, SHARC has dominated the floating-point DSP market, delivering exceptional core and memory performance complemented by outstanding I/O throughput. At 319 MFLOPS per dollar, the SHARC family makes floating point more accessible than ever. And for high performance multiprocessing, TigerSHARC Processors deliver GFLOPS performance with glueless, 1 Gbyte/sec multiprocessing link ports, and 24 Mbits on-chip memory.

If you're ready to speed your time to market and showcase your design's real potential, Analog Devices SHARC and TigerSHARC Processors deliver floating-point performance—and performance density—like never before.

© 2005 Analog Devices, Inc. Product and company names listed are trademarks or trade names of their respective owners.

SHARC Processor Leadership in MFLOPS/$

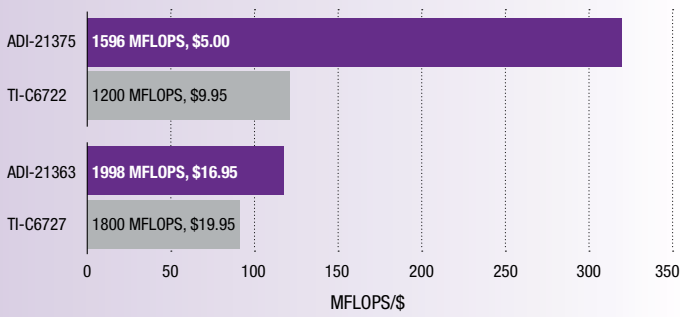

Leading Performance vs. Competition

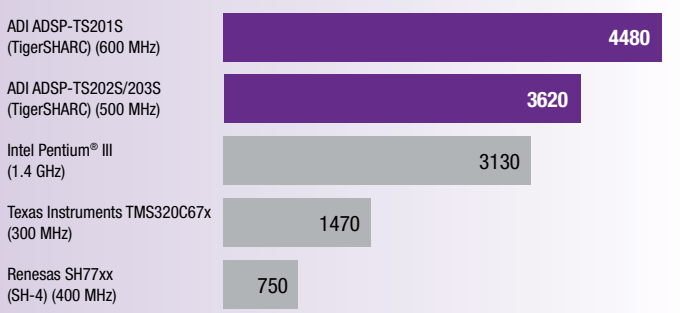

BDTi BENCHMARK PARTNER

The BDTImark2005™ provides a summary measure of DSP speed. For more information and scores see www.BDTI.com. Scores© 2005 BDTI.

"Floating-point processing essentially eliminates concerns about dynamic range and ensures precision in the most intensive and demanding signal processing tasks—but it's been an expensive alternative in terms of both cost and power. Architectural advances are now making this capability far more widely accessible."

— Dr. Alan Oppenheim
Ford Professor of Engineering
Massachusetts Institute of Technology

SHARC Advantages

- Broad portfolio of code- and pin-compatible processors starting at $5 and providing 2.4 GFLOPS of signal processing performance
- High precision computation enabled by SHARC's 32-bit fixed- and floating-point architecture
- Up to 400 MHz core instruction rate
- Single-Instruction Multiple-Data (SIMD) computation architecture
- Simplified system design and lower system cost via a wide array of integrated peripherals

TigerSHARC Advantages

- Optimal balance of computational performance, large on-chip memory, and I/O bandwidth maximizes sustainable FFT performance
- Memory integration up to 24 Mbits accessible at core frequency enables internal execution of FFTs of 64 K points maximum
- Glueless, scalable multiprocessing capability with 5 Gbytes/sec aggregate throughput
- Interlocked pipeline for "out-of-the-box" execution of ported workstation code
- 24 Mbits of on-chip RAM—three times more than the closest competition

Mackie TT24™ Digital Live Console
Mackie uses SHARC's on-board DSP power to provide audio dynamics, equalization, and effects processing in a way that's practical, intuitive, and fast enough for live sound.

BMW M5 Series
BMW chose SHARC Processors for control of its sophisticated 500 hp V-10 motor. The V-10 engine control unit performs an amazing 200 million calculations per second to ensure optimal performance.

R&S® Radio Communication Tester CMU
With ADI floating-point processors, Rohde & Schwarz provides a versatile mobile radio test platform that conforms to multiple international standards and delivers maximum measurement speed for extensive testing of current and future mobile radio networks.

iCyt™ Cell Separation Instrument
TigerSHARC Processors are used in a high throughput, biological cell separation instrument designed and built by iCyt Visionary Bioscience, Inc. while utilizing Bittware's TigerSHARC development boards.

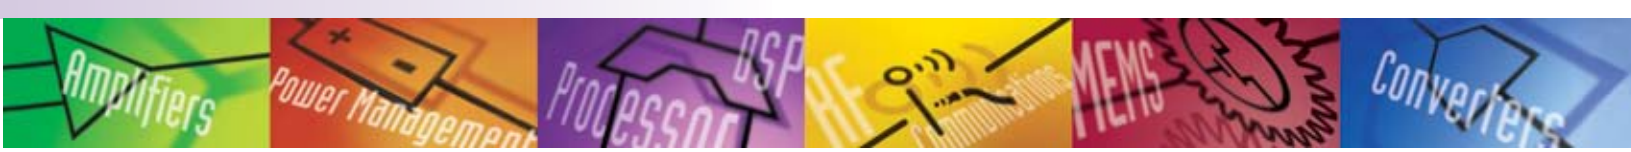

BittWare

Visit **www.analog.com/perform** for information on special bundled board offers from BittWare, Inc. and Analog Devices.

ANALOG DEVICES

High Availability and Interoperability for Telco Class Applications

Ensuring performance standards are met using an ATCA application-ready platform can offer reduced time-to-market and maximize interoperability via industry standards. Here is one approach to validating interoperability and performance.

by Jeff Hudgins and Cynthia Tucker
Alliance Systems

Interoperability and performance are key market enablers for network equipment providers. The relationship between performance and time-to-market has been widely studied and found to be strongly correlated; however, the performance gap that exists between an interdependent (proprietary) design and a modular (standards-based) design is not as well explored.

Implementing platforms using proprietary hardware and middleware involves a higher level of investment in development costs and slower time-to-market. Ideally, open modular standards like ATCA should ensure "interoperability" between multiple hardware and middleware vendors. Making everything work to ensure performance and interoperability is expensive and takes considerable skill.

Equipment providers are under constant pressure to reduce costs, add new functionality and shorten the time it takes to deliver new platforms. As the adoption of ATCA increases, the need to evaluate and adapt to newer technologies presents interesting challenges. To address this challenge, Alliance conducted a series of tests over a four month span on our A-Series 5000 ATCA platform to measure performance and interoperability within multiple configurations required by OEMs and application developers.

Basic ATCA Platform Building Blocks

ATCA systems comprise many of the components that enable functionality for wireless edge and access applications, core transport, and media gateways and servers. The building blocks of a

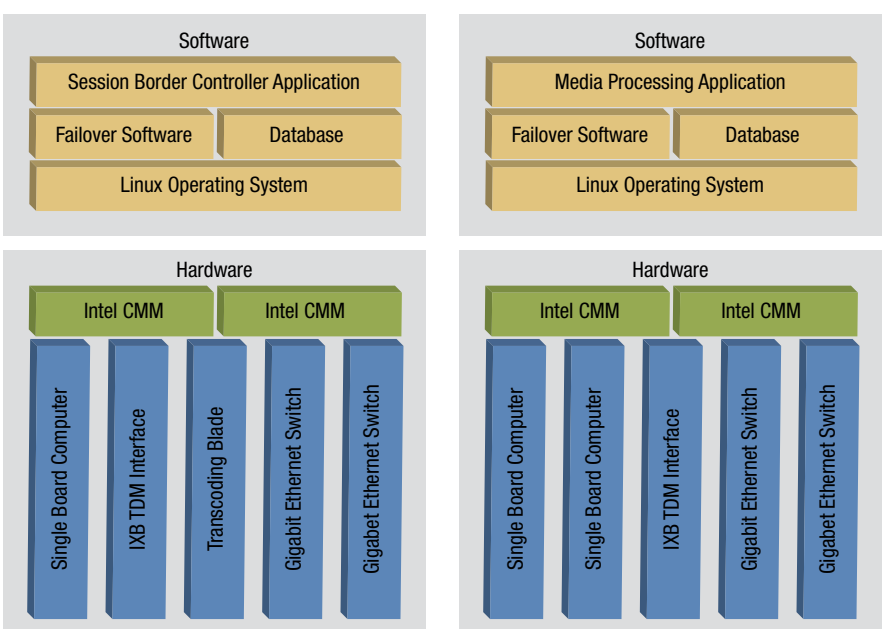

Figure 1 Architectural layers within an ATCA ARP.

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

SHARE YOUR VISION.

WE'LL SHARE OUR INNOVATION.

Share your vision with us. We'll customize our products to your requirements or partner with you to develop custom products all the way through deployment. Either way, you'll get leading-edge board level solutions for the most demanding embedded applications. Tell us what you need at **info@vadatech.com**.

- Comprehensive Hardware/Software solution for Intelligent Peripheral Management Interface (IPMI) version 2.0
- Symmetric Multi-Processing CPU for AMC/VME Modules
- A complete line of ATCA and AMC carriers
- High dynamic range A/D Converters

vadatech inc.
THE POWER OF VISION

www.vadatech.com 702.896.3337

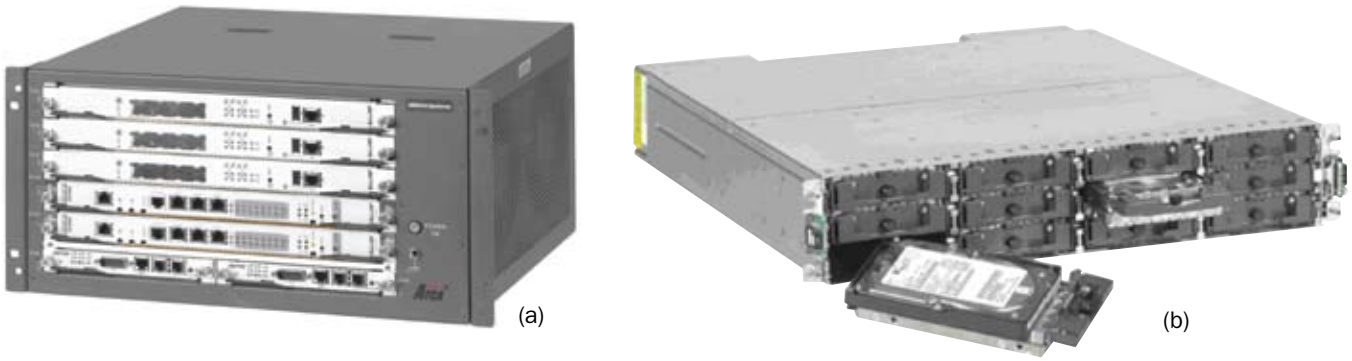

Figure 2 Alliance's A-5000 ATCA platform (a) and NS-2500 NEBS-certified FC storage system (b).

complete ATCA *application-ready platform* (ARP) also include integration and testing of many hardware and software components to enable deployment of triple and quadruple play applications. The fundamental building blocks as shown in Figure 1 include platform hardware, media processing and IXB boards, Intel chassis management modules (CMMs) and processing boards, failover software, database software and operating systems, all of which must meet carrier-grade, high-reliability requirements such as NEBS, ETSI and a minimum of five nines (99.999%) availability.

Compliant with the PICMG 3.0 and 3.1 specifications, the system selected for our series of interoperability and performance tests is Alliance's smaller form-factor A-5000 ATCA 5U rackmount NEBS-compliant server shown in Figure 2a. The server can achieve levels of backplane interconnect bandwidth up to 3.125 Gbits/s. Standard configuration includes two Intel Chassis Management Modules (CMM) and two 16-port PICMG 3.0 Gigabit Ethernet base fabric switches for redundancy and provides three node slots for expansion boards such as low-voltage single board computers supporting dual Xeon 2.0 GHz processors, media server, Intel IXP network processing and other ATCA-compliant blades.

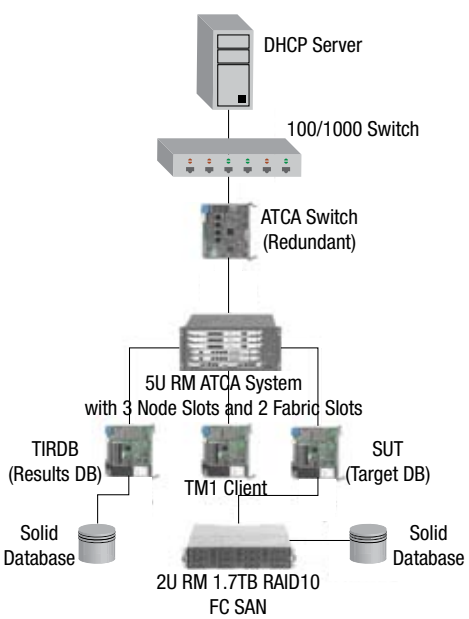

Figure 3 ATCA Test Platform architecture with storage system

The test platform included CPU blades (one blade with a dual-port 2 Gbyte FC HBA), GbE switches and two Intel shelf managers. It also included dual Intel 2.0 GHz Xeon processors for the CPU blades with an Intel E7501 chipset supporting a 533 MHz FSB and dual GbE network interface cards routed to backplane. Each Intel CPU blade contains a 60 Gbyte enhanced duty ATA/100 2.5-inch hard drive and either 1 or 2 Gbytes of ECC Registered PC2100 DDR RAM.

The NS-2500 2U rackmount NEBS-certified, Fibre Channel (FC) storage system was introduced into the test environment to validate performance and interoperability levels within an external storage environment. The SAN was configured with twelve 300 Gbyte FC drives in a 1.7 TB dual-channel RAID10 (Figure 2b).

Validating Interoperability and Performance

Platform qualification and validation testing to support high-reliability application environments is extremely time-consuming and costly.

To address both time and cost, we selected the Design of Experiments (DOE) approach to analyze a variety of variables and outcomes. Analysis of the Design of Experiments is a collection of models in which the observed variance is partitioned into components due to different factors estimated and/or tested as compared to one factor at a time approach. This approach was selected over a "one change at a time" testing method to uncover the effects of any interactions and/or dependencies and can validate results in one-fourth the time. The chart in Table 1 shows the input variables and levels used during the analysis.

IOMeter and Telecom One, industry standard benchmarks, were used to measure performance factors.

- IOMeter is a random access time test. It varies requested block sizes as well as issues new requests before a previous request is fully serviced thus driving up the simulated load.
- Telecom One (TM1) is a database benchmark test developed by Solid (www.solidtech.com, Cupertino, CA) designed to simulate a typical Home Location Register (HLR) database used by mobile carriers. Results show mean qualified throughput (MQTh) of the simulated HLR database, and the response time distributions for seven predefined transactions that insert, update, delete and query the data with differing frequencies.

To determine I/O throughput, IOMeter test script configurations were used on the local partition-less hard drive and the external Fibre Channel storage array. Using a 2 Mbyte file, the test ran for ten minutes on the 2.0 GHz and 2.8 GHz blade within both 1 and 2 Gbyte memory configurations.

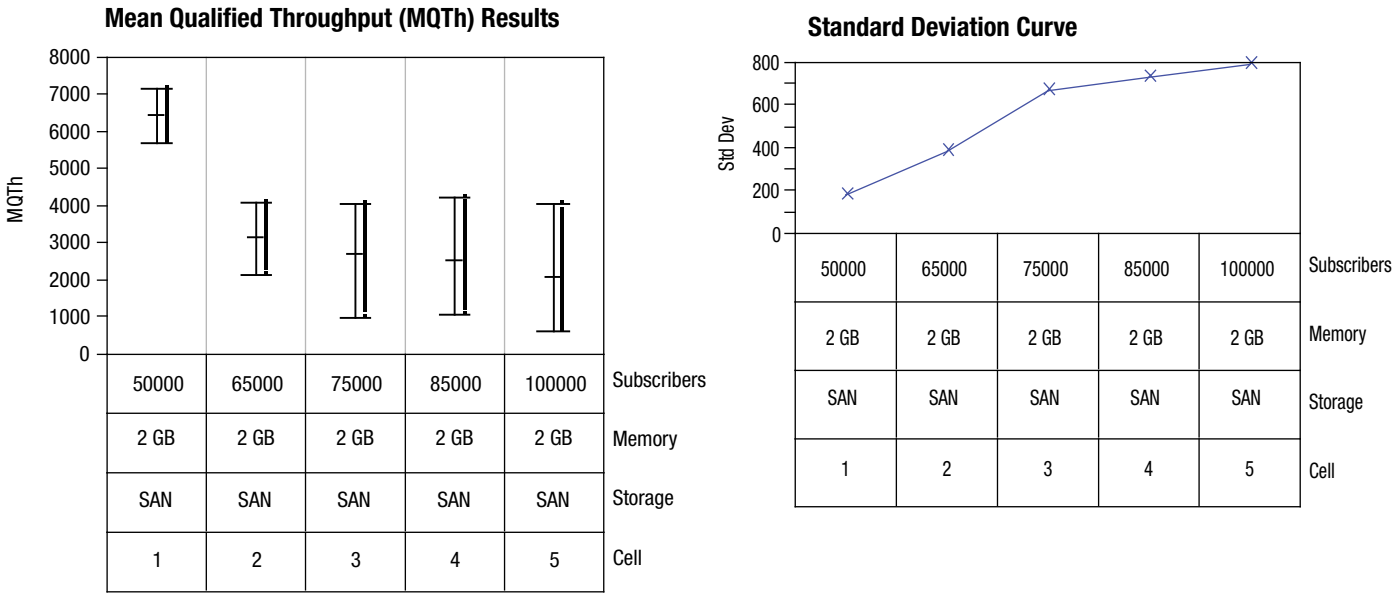

Figure 4 Telecom One HLR Benchmark Tests (65,000, 75,000 and 85,000 subscribers).

IOMeter test results indicate that storage played a major factor relative to I/O performance—summary statistics in the 1 Gbyte memory configurations with the 2.0 GHz blade using SAN storage produced transfer rates of 83.3 and 50.6 Mbytes/s (read/write) versus 30.2 and 30.5 Mbytes/s (read/write) on the local drive.

Further random access performance gains were achieved using the external SAN storage on the 2.8 GHz blade—transfer rates of 132.7 and 68.0 Mbytes/s (read/write) respectively. And similar to the 2.0 GHz blade results, we found no significant performance improvements in the 2 Gbyte memory configuration versus the 1 Gbyte memory configuration.

The TM1 HLR benchmark test environment shown in Figure 3 included the ATCA server with three CPU blades to the DHCP server in a closed network environment. Each compute blade had the simulated HLR database installed—one blade (the CPU blade with the FC HBA) was designated as target database, one blade was designated as the database results server (TIRDB) and one as the TM1 benchmark client. Each blade was running Red Hat Enterprise Linux version 3 update 5.

Test results indicate the addition of external SAN storage had a significant impact on throughput at the 50,000 subscriber level. Throughput variations increased significantly as the subscriber count increased to 100,000. And as in the previous test results, increased system memory had no effect relative to throughput during our tests.

Input Variables	Levels
Number of Subscribers	50,000 65,000 75,000 85,000 100,000
Memory	1GB PC2100 2GB PC2100
Storage	Local 60GB External 3.6TB Fibre Channel
Processor	Dual 2.0GHz 400MHz FSB Dual 2.8GHz 800MHz FSB

Table 1 Input variables and associated levels for DOE analysis.

Subscriber levels were also defined at 65,000, 75,000 and 85,000 using the external storage array. Benchmark results indicate that as subscriber counts increase to 65,000, throughput declines rapidly. Less significant decreases at 75,000 subscribers are shown in the HLR benchmark charts (Figure 5). CPU utilization did not appear to be a bottleneck as subscriber levels increased to 100,000. In subsequent tests, however, the impact of using the Intel Nacona chipset on the Adlink 6890 provided improved I/O performance.

Throughout the testing, Alliance was able to validate interoperability of the A-5000 with different vendor elements, and benchmark performance characteristics of operations within both local and external storage systems. This level of validation testing and systems integration provides the framework and building blocks to transition OEMs from proprietary environments to ATCA standards.

Results indicate that additional compute power would be required to maintain performance levels. This was validated in subsequence tests as MQTh improved with the introduction of dual Intel 2.8 GHz Xeon processors on the Adlink CPU blade with an Intel E7520 chipset supporting a 800 MHz FSB. ◢

Looking For More?
Visit www.rtcmagazine.com to download additional technical information related to this article.

Alliance Systems
Plano, TX.
(972) 633-3400.
[www.alliancesystems.com].

A Chip Off The Old Block.

Good genes run in our family. Ampro's EPIC and Mini-ITX embedded motherboards have inherited many characteristics from their PC/104 and EBX cousins. These low cost SBCs offer many of the embedded features found in Ampro's rugged CoreModule™, LittleBoard™ and ETX products. Features such as fanless operation, serial ports with RS422/485 options, dual Ethernet, battery-free operation, USB and LAN boot and a host of others. Plus, Ampro's industry-leading lifecycle management policies apply, assuring you of long lifecycle and consistent product deliveries, time after time, with notification of product changes well in advance.

Don't let the low price fool you. If you're looking for an embedded motherboard for a cost-sensitive application, the best solutions are Ampro ReadyBoard™ and MightyBoard™ SBCs. With a lineage second to none, our embedded motherboards have Ampro written all over them.

Quality...
Durability...
Dependability...
It runs in the family.

Ampro Embedded Motherboards

Call today (408) 360-0200 or visit: www.ampro.com

AMPRO®

Ampro is a registered trademark of Ampro Computers, Inc. All other trademarks are the property of their respective owners. © 2005 Ampro Computers, Inc. All rights reserved.

ATCA300: Expanding to New Applications

Specifically designed to meet the needs of transport equipment makers, the emerging ATCA300 specification leverages off of ATCA 3.0 and is particularly of interest for fiber-optics communication.

by Niall Robinson
Optovia

Most major telecommunications equipment manufacturers were quick to embrace AdvancedTCA (ATCA) as the basis for a standards-based, modular platform for blade/media servers in wireline and wireless networks. This is because following the ATCA standard decreases the time, money and resources that they would otherwise expend designing, developing and creating portions of their products that are common to all manufacturers.

Use of this open standard platform has promoted best-in-class system design and reduced market- and product-development risk for all companies involved in the telecom equipment supply chain. Carriers are enjoying faster delivery of new service-enabling applications and solutions from their equipment suppliers. Subsystem vendors also find they are able to support more customers and leverage their products' versatility in multiple markets.

New Horizons

Realizing all the benefits that telecom manufacturers gleaned from the ATCA 3.0 standard, transport equipment manufacturers decided in late 2004 to try to duplicate its success in the $8 billion transport equipment market. Intel, Marconi, Optovia, Schroff and Bustronic stepped up to sponsor a PICMG subcommittee to develop AdvancedTCA300 (ATCA300) for high-end, highly reliable and high-port-count transport equipment.

Today, a total of 56 companies are registered members of the subcommittee, which is working to develop the necessary requirements for a platform that accommodates the unique form-factor and specifications of transport equipment. To date, ATCA300 has reached its second draft with completion expected early in 2006. The ATCA300 subcommittee has endeavored to make as few changes to ATCA 3.0 as possible. As a result, ATCA300's backplane, shelf management architecture and backplane power architecture are the same as ATCA 3.0.

Like ATCA 3.0, ATCA300's common software architecture enables transport equipment makers to significantly reduce both the cost and the time it takes to develop software for their equipment. Because software is the costliest and most time-consuming component of transport products, ATCA30's0 common software architecture speeds time-to-market and fosters an ecosystem that streamlines the transport equipment supply chain overall.

Although changes have been kept to a minimum, some have been necessary. First, to meet the needs of the harsher transport environment, ATCA300 adheres to requirements for telecommunications gear outlined by NEBS and ETSI standards. ATCA300 also has a different depth dimension. As its name suggests, the platform enables a shelf structure that supports 300 mm deep equipment practices. And, the PCB that plugs in from the front was shortened from 280 mm in ATCA 3.0 to 220 mm for ATCA300.

Perhaps the most striking difference between the two is that the ATCA 3.0 rear termination module (RTM) was eliminated. This is because all transport equipment must be capable of being flush-mounted either against the wall of a hut or central office (CO), or back-to-back with another box when deployed in the center aisle of a hut or CO.

The remote network sites, or huts, in which transport gear is typically deployed, most often have four "lanes" worth of equipment. Two lanes are mounted against each wall and the other two are placed back-to-back down a center aisle. There is only 5 mm to 10 mm behind each box, which leaves no room for maintenance or access from the rear. Access to ATCA300 for any purpose, including maintenance, is provided only from the

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

front. Even though space is not as big an issue in the CO, transport equipment used in the CO still adheres to the requirements noted above.

In lieu of the RTM, the ATCA300 offers a front termination module (FTM), which is an extension to the chassis. This secondary shelf is connected to the backplane of the primary shelf (Figure 1). The ATCA300's backplane is divided into several zones:

- Zone 1 is used for power.
- Zone 2 contains the management and traffic interface.
- Zone 3 can be used to communicate with, and connect to, the FTM, which defines a Zone 4 backplane.

Best in Breed

With the advent of ATCA300, the days when transport equipment's proprietary physical footprint was enough to secure market share, are over. ATCA300 levels the playing field for the vast array of optical technology companies that will now compete with one another exclusively on the competence of their technology. In effect, ATCA300 is the new Multi-Source Agreement for the next generation of transport products.

As with ATCA 3.0 standard devices, no one actually manufactures "plain vanilla" ATCA300 equipment. Manufacturers develop their own equipment to conform

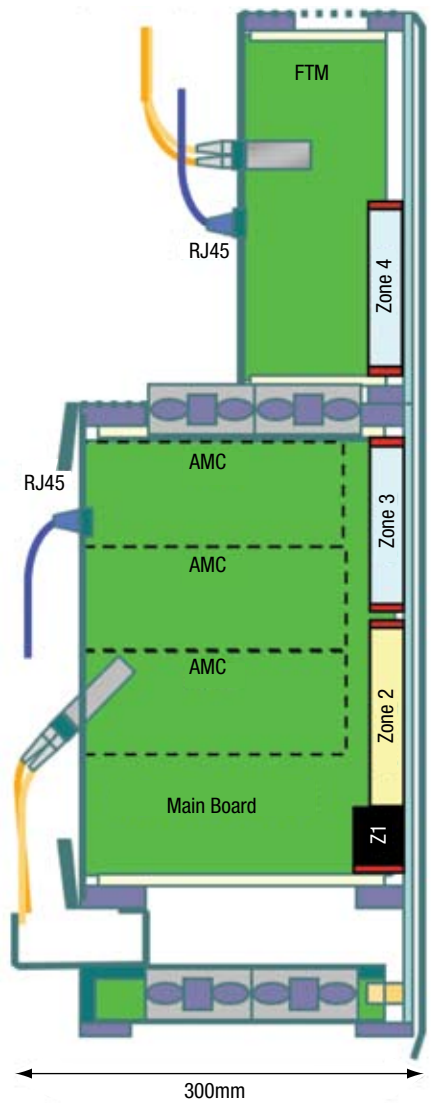

Figure 1 Side view of ATCA300 shelf showing main board and FTM. The ATCA300 standard seeks to specify what Zone 4 looks like and how it connects to, and is managed by, the main shelf. Smaller than the main front board, the FTM is only 110 mm deep, which allows for additional cabling plenum to handle significant I/O capability.

Is your cPCI system going end of life?

Look no further... Hybricon has the solutions !

Hybricon® Innovation • Quality • Service

Bridging the gap between legacy products and you..

- ✓ 3U high 5 Slots 6U x 160 mm cPCI
- ✓ N+1 redundant hot swap power supplies
- ✓ Slim line drive shuttle
- ✓ Hot swap fan tray
- ✓ Cools up to 60W /slot
- ✓ Front to rear cooling
- ✓ Alarm panel (optional)
- ✓ Designed to meet NEBS level 3
- ✓ Available in AC or DC versions
- ✓ RoHS - 6 compliant

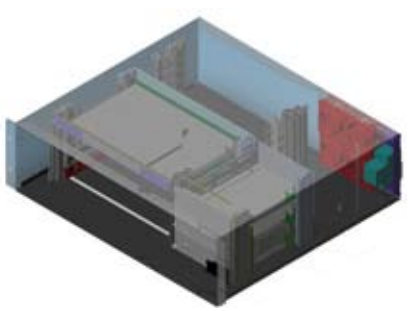

www.hybricon.com

Hybricon Corporation 12 Willow Road Ayer, MA 01432
ISO 9001-2000 Certified

Call Today
1-877-HYBRICON

Figure 2 Example of a carrier-class, highly available, five-slot chassis (C4000 from Optovia)—a pre-standard, 300 mm-compliant chassis with an ATCA backplane. The multiple redundant and field replaceable units include fans, shelf processors and power entry modules. The device complies with NEBS Level 3, GR-1089-Core, ETS300-019, ETS 300-386, and UL/CSA EN60950, CB-TUV standards and specifications.

to PICMG specifications. The equipment can vary, but only in a limited manner. For instance, boxes can accommodate horizontally or vertically mounted shelves.

There is also room for creativity in the way ATCA300 is deployed. For instance, a five-shelf horizontal chassis can be cascaded to another two virtual shelves to enable customers to scale their networks. In short, just how ATCA300 boxes actually come together is left to market innovation. One thing all ATCA300 boxes do have in common is that they enable any blade to be plugged into any chassis. Eventually, they will be able to manage multiple blades across multiple vendors.

Forging Ahead

Even though the market is in a pre-standard position today, ATCA300's inroad into the transport segment has already begun. Critical elements of the transport network such as amplifiers, 10 Gbit/s transceivers and reconfigurable add drop multiplexers (ROADMs) are already available or in development programs that will have them ready for release this year. Products currently under consideration include an MSTP and routing platforms, intelligent optical multiplexers/demultiplexers and 40 Gbit/s terminal equipment.

During the next two years, the ATCA300 open platform ecosystem will become populated with components and parts of the end product that previously would have had to be designed and manufactured during the entire life cycle of product development (Figure 2). As a result, the ATCA300 environment will foster a strong supply chain as well as high volumes that will lower unit costs for all parties in that chain.

Demonstrations of an ATCA300-based transport system are slated to take place at the Optical Fiber Communication Conference, OFC 2006, in Anaheim, CA, March 7-9. In the meantime, additional development continues on a front panel extender, which will enable the ATCA300 blade to fit inside an ATCA 3.0 chassis. This full backward compatibility will reduce the risk of ATCA300 development by enabling ATCA300 blades to take advantage of the installed base of PICMG 3.0 equipment.

Looking For More?
Visit www.rtcmagazine.com to download additional technical information related to this article.

Optovia
Acton, MA.
(978) 266-2830.
[www.optovia.com].

ELMA SYSTEM INTEGRATION
The Final Piece of Your Embedded Packaging Solution
At last, a complete embedded packaging solution that's simple and reliable! You'll work with an experienced engineer who will see your project through design, prototype, production and integration. Utilize our testing and verification resources to make sure your application meets agency standards, and your specifications. With everything done in one place, you save time and money on logistics. All you have to do is sit back and watch all the pieces fall into place.
ELMA
Your Solution Partner
USA Elma Electronic Inc.
Phone: 510.656.0606 Fax: 510.656.8008 E-Mail: sales@elma.com Web: www.elma.com
© 2003 Elma Electronic Inc.
Standard Enclosures
cPCI/2.16, ATCA, VXS, VME/64X, Custom
19" and 23" Rack mount
1U to 14U
2 to 21 slots
IEEE 1101.10/.11
STMP
NPI Solutions
Design
Prototype
Testing
Manufacturing
Assembly
Certification
Integration

All the Linux tools you need in one spot.

NexusWare™ is a suite of Linux-based software products that enables users of our application platforms and blade products to rapidly develop and deploy value-added capabilities with their solutions.

NexusWare Core
Linux-based development, integration and management environment—the foundation for the NexusWare installable software packages below:

NexusWare WAN
WAN protocol support for X.25, Frame Relay, HDLC, Radar Receiver and more.

NexusWare C7
MTP-2 starting point for building SS7 applications.

NexusWare MG
Protocols and tools for media gateway control, IP call control and PSTN signaling.

NexusWare ISM
Programming and control tools for system-level intelligent shelf management.

Learn about our NexusWare software at:
www.pt.com/linux3

Platforms. Components. Software. Support.

Many solutions. One vendor.

PERFORMANCE TECHNOLOGIES

CELEBRATING 25 YEARS
PERFORMANCE TECHNOLOGIES

Communications Processors vs. Network Processors: Programmable Data Plane Approaches

Network processors and communications processors are used very differently for implementing programmable data plane functionality in communications processing. Designers must understand these differences to ensure that their designs will perform the required functions with the ability to meet changing and expanding requirements, within a board that meets the required physical characteristics.

by David Smith
Freescale Semiconductor

Many embedded systems, such as routers, data/voice switches and wireless base stations, require some level of communications processing. Typically, communications processing is split between the data plane, which manages the data flow, and the control plane, which manages communications (Figure 1).

Two different silicon design approaches are available for implementing programmable data plane functionality: network processors (NPs), optimized for packet processing throughput, and communications processors, which integrate data and control plane functions and other peripherals into one system-on-chip. These two different approaches are used very differently by designers to implement data plane functions, as well as within a system to perform the overall system application.

The key characteristics of a communications processor not available from a network processor include a high "real performance" level of the general-purpose core; highly efficient use of the general-purpose core in data plane and control plane tasks; a single, shared memory space for data plane-specific processing and general-purpose processor hardware; and simplified hardware implementation that provides less power, space and cost than an NP-based solution.

Data Plane Functionality

A communications application converts and forwards serial traffic—streams, cells or packets—from one communications port to another. Data plane functions are critical to the traffic flow rate. The control plane comprises maintenance functions not directly required to forwarding

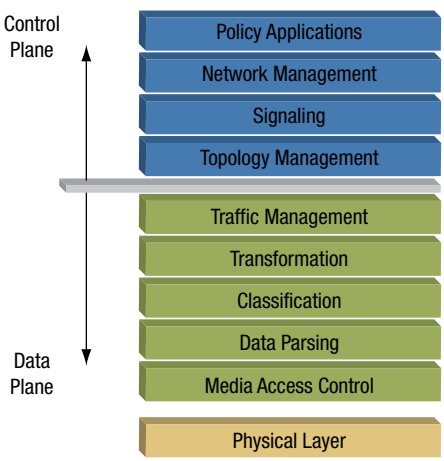

Figure 1 Communications processing functions are divided into functions required on the data flow (data plane) and functions required to manage the communications (control plane). Data plane functions are critical to the traffic flow rate, while control plane functions are maintenance functions not directly required to forward traffic through the system.

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

PCI Express Connectivity
10Gb per Second
Imagine The Possibilities...

MAX EXPRESS™

One Stop Systems provides complete PCI Express expansion solutions for motherboard, SHB (System Host Board), and CompactPCI Express architectures.

ONE STOP SYSTEMS
2235 Enterprise, Suite 110
Escondido, CA 92029
Tel (760) 745-9883 Fax (760) 745-9824

Call today.
(877) GET-CPCI
www.onestopsystems.com

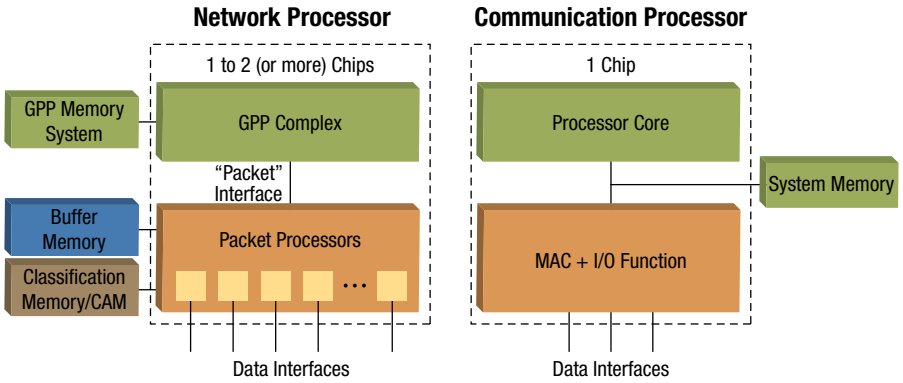

Figure 2 A network processor-based architecture often requires more chips compared to a communications processor design, resulting in loss of valuable space and increased power requirements and a higher cost per board.

traffic through the system. Control plane functions can usually be performed at a lower priority than data plane functions.

The data plane's several layers have different functions. Media access control (MAC) controls wire/fiber/media access to the physical network. It performs serial-to-parallel conversion, finds and maintains media synchronization and transfers data traffic into system memory.

The data parsing layer separates forwarding or classification information from the data unit and passes it to the classification function. This layer determines what needs to be done with incoming information. This includes determining the egress port, whether the data is management control information, priority and filtering.

The transformation layer converts the received data unit into one that can be sent on the media associated with the egress port. This often includes an update of source/destination fields and CRC calculation. More extensive transformation may be required if the data is to be forwarded between ports with dissimilar media interfaces.

Traffic management is used to implement a priority or congestion management scheme. Topology management builds and maintains information on overall network connections and mapping, and ensures that lookup tables used by the classification function are created and maintained.

The signaling layer performs higher layer connection management and recovery.

Network management builds and communicates information for overall network status and health. The policy applications layer provides additional network information and control, such as provisioning, billing and security.

Preferred Traits of Programmable Data Plane Solutions

There are numerous, distinctly different silicon approaches to implementing data plane functionality, ranging from hardware- to software-centric. Hardware-centric approaches include ASICs, FPGAs and ASSPs, which trade off reduced flexibility for maximum performance and/or minimum cost. These approaches are rarely used for embedded system designs because they are designed for specific system applications; embedded board vendors typically do not offer hardware-centric solutions.

However, numerous programmable silicon board-level solutions are available, enabling equipment vendors to bring differentiated products to market with a reasonable amount of skill level and development resources. The primary processor choices for these board solutions are network processors and communications processors with embedded general-purpose processor cores.

Of these two choices, communications processor-based solutions have gained a higher market acceptance for several reasons. These include a standard, general-purpose processor programming model, ease in mixing control and data plane functionality, flexibility in application mapping and simpler, more efficient hardware design implementation.

Since communications processors

ACROMAG I/O SOLUTIONS

ANALOG I/O DIGITAL I/O SERIAL I/O FPGA COUNTER/TIMER QUADRATURE

Industry Pack I/O Modules

If you need a variety of I/O functions or have a slew of signals to process, then Acromag Industry Packs are your solution. Plug in assorted I/O modules on a carrier card and you can replace several boards. Or, insert multiple IPs of the same type for a single-slot solution with tremendous capacity. Now *that's* flexibility.

Acromag offers the widest selection.
Our Industry Pack modules and carrier cards all come with an unbeatable combination of features, performance, and value.

- Choose from more than 100 IP modules
- Carriers for 3U/6U VME, CompactPCI®, PCI
- Software drivers for RTOS and Windows®

Flexibility means options. And Acromag has more embedded I/O options than anyone.

For flexible, capable, and reliable I/O solutions, think Acromag.

Acromag®
THE LEADER IN INDUSTRIAL I/O

www.acromag.com
800-881-0268 or 248-624-1541

Embedded Systems Conference

Silicon Valley Conference: April 3-7 Exhibition: April 4-6

06

FIVE DAYS.
ONE LOCATION.
ENDLESS POSSIBILITIES.

Early Bird Registration
Register by March 7 and Save $200
on Conference Packages
www.embedded.com/esc/sv/registration
http://www.embedded.com/esc/sv/registration

McEnery Convention Center, San Jose, CA

www.embedded.com/esc/sv

CMP
United Business Media

Sponsored by: EETIMES Embedded Systems Design

Embedded Systems Conference Silicon Valley

use a standard, general-purpose processor programming model, designers can use standard programming languages running under standard operating systems, as well as a standard development environment, debug/simulation tools and OS specifics.

A network processor approach, however, requires engineers to relearn many software skills in order to use the data-specific processing resource, increasing time-to-market. In many cases, network processor vendors must perform the programming for equipment vendors. Equipment designers should be leery of this approach, because it does nothing to grow the skill level of their engineering talent, and it reduces the opportunity for product differentiation, since network processor vendor code will be available to all users.

Communications processors provide ease in mixing control and data plane functionality. By contrast, a network processor approach forces a rigid architectural division of labor between the hardware elements that perform control and data plane functions. In many cases, these functions are intertwined, and the network processor approach can force an unnatural division of control/data plane functionality that reduces efficiency.

Since the general-purpose core within a communications processor typically performs control plane functions and is capable of data plane functions, it is much easier to architect a solution requiring tight coupling of control and data plane functions. Exception-path handling is a prime example. Because of the relatively small code space available to packet processors, most network processor-based applications pass this on to the control processor. With a communications processor-based solution, exception-path functionality is easily integrated into "normal" path code routines.

In all network processor implementations, multiple packet processors are required to meet performance goals. This architecture dictates that the entire data plane functionality be split among the packet processors. As product functionality requirements change, significant remapping of these resources may be required to accommodate the new function. This is especially true for pipelined packet processor architectures. Examples of functional changes include expanding

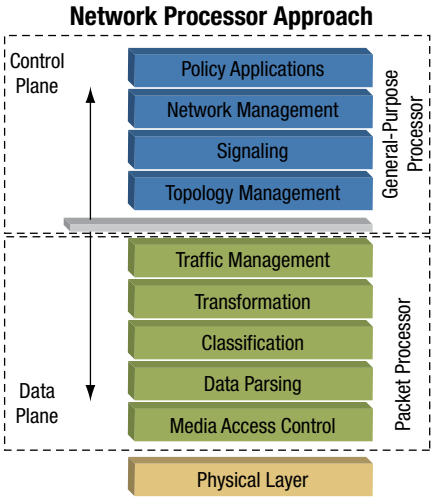

Figure 3 A network processor may include packet processors intended to perform data plane functions; in reality these may not perform efficiently.

a product to support a new protocol and adding capability to handle system exception and recovery conditions that were not understood until deep into the product design cycle.

Communications processors do not require intensive resource mapping. Adding new functionality to a general-purpose processor involves simply writing a new software function with little to no impact on existing code. Even error and excep-

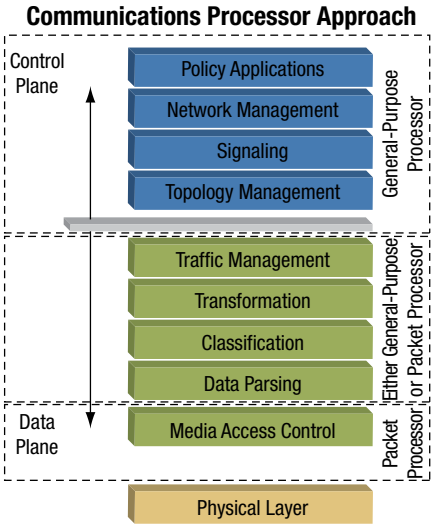

Figure 4 A communications processor provides extensive flexibility as well as ease-of-programming by a skilled base of software engineers.

ACROMAG I/O SOLUTIONS

ANALOG I/O DIGITAL I/O SERIAL I/O FPGA COUNTER/TIMER QUADRATURE

PMC I/O Modules

Looking for higher performance without sacrificing the flexibility of a mezzanine format? Acromag PMCs enable faster data transfer rates with all the benefits of a PCI interface.

Acromag's wide variety of I/O modules includes high-speed A/D, high-resolution D/A, high-density digital I/O, and even extended temperature ranges. Now *that's* capability.

New FPGA I/O means greater capability. Re-configurable FPGAs let you design custom I/O boards much faster and at a fraction of the cost. Download your unique code into the FPGA for unlimited signal processing possibilities.

- Virtex-II® FPGA with up to 3 million gates
- 1MB on-chip RAM, 9MB on-board SRAM
- Variety of analog & digital I/O configurations

For flexible, capable, and reliable I/O boards, think Acromag.

Acromag®
THE LEADER IN INDUSTRIAL I/O

www.acromag.com
800-881-0268 or 248-624-1541

Master Multicore Hardware & Software Designs!

Register by March 10th and Save! Full Conference only $495 1 Day Pass only $295

multicoreEXPO™

The Fast Track to Multicore Solutions

March 21-23, 2006, Santa Clara Convention Center

Join the industry's top experts in a unique conference that delivers real developer solutions for multicore and multiprocessing designs.

In the 3 days of this conference, delegates will learn about:

- New multicore hardware and software development techniques
- Hot products to facilitate your development of multicore-enhanced applications
- Practical, hands-on approaches to designing multicore solutions
- Key developments and roadmap of the Multicore Association.

Event Sponsors: (As of January 20)

Platinum PLUS

Sun microsystems

Platinum

ARC, picoChip, WIND RIVER

Gold

ARM

Silver

ALTERA, codito, Green Hills Software, red PLAIN, tensilica, American Arium, Evidence, imec, eL express logic, Teja, MIPS Technologies, PolyCore Software, TILERA, LYNUXWORKS™

Media Sponsors:

EETimes, IQ, Circuit Cellar, RTC, COTS Journal

Sponsored by the:

Multicore Association

Don't miss the Multicore Expo! Registration opens February 1.

www.multicore-expo.com

tion conditions can usually be handled by adding code to detect the condition in the existing code and then writing a modular, "self-contained" routine.

Many network processor implementations require external assistance, such as classification engines, traffic managers and media access controllers. Within a communications processor, these functions are either performed by the general-purpose core or integrated function blocks. Within a network processor-based system, there is a rigid split between packet and general-purpose processors (Figure 2). At a minimum, all practical network processor-based solutions require a separate general-purpose processor to handle the control plane and possibly some exception-path processing. Because network processors are architected in a "wire rate or bust" fashion, extra memory in the form of buffer memory, usually SDRAM, and classification-specific memory, SRAM or CAM, is required. This is in addition to the memory required for the general-purpose control processor.

Communications processors require only a single system memory space. By efficient use of caching, the functions requiring low latency access can be accommodated without the added complexity of multiple memory systems. Because of the requirement for multiple, disparate memory systems and the likely need for "bolt on" assistance to perform tasks that are too complex to be performed by packet processors, a network processor-based design usually has a very high chip count compared to a communications processor-based approach. The extra chips consume valuable board space, power and cost.

Distinguishing Network Processor from Communications Processor Approaches

Because of the limited market adoption of network processors, some silicon vendors are introducing solutions that are network processor-based, but that may be marketed as communications processors. Both types of devices may have a built-in general-purpose processor core along with one or more integrated packet processors. However, the approach is still fundamentally a network processor approach if the general-purpose processor cannot be efficiently used to perform data plane functions.

In a network processor approach, the general-purpose processor primarily performs control plane functions. A communications processor approach may also have a low-level, on-chip packet-processing capability, but this may or may not perform all the data plane functions. Instead, the part is architected such that the general-purpose processor may efficiently perform much of the data plane functionality. The key architectural difference is that a network processor requires a pseudo-packet memory-to-memory transfer within the system (Figure 3). However, a communications processor, with its single memory system, does not require this memory transfer step (Figure 4).

Extracting Maximum Data Plane Performance from General-Purpose Processors

Communications processors have been an industry mainstay for 20 years and have been used for data plane processing applications, such as routers, voice gateways, wireless infrastructure (gateways and base stations), DSLAMs and telephony equipment. Leveraging the exceptional flexibility of communications processors, equipment vendors can implement innovative equipment solutions, often defining entirely new product and market segments. Over time, a wealth of collective experience in real-world designs has shaped the evolution of communications processors to meet the data plane performance requirements of communications system designers.

One major design technique is the use of caches. Modern general-purpose processors have at least one level of caching. Cache allows the code and data critical to data path forwarding to be accessed quickly by the core, reducing the time the core spends processing the traffic. Larger caches allow more code and data to be local to the core. Some cache architectures allow certain areas of memory to be locked into cache. Typically, this feature is enabled by running the "normal" path of data plane code to load the cache and then by applying the lock. More locking options and smaller granularity of lockable sections provide the data plane code designer greater flexibility.

ACROMAG I/O SOLUTIONS

ANALOG I/O DIGITAL I/O SERIAL I/O FPGA COUNTER/TIMER QUADRATURE

PCI and CompactPCI® I/O

Designed for industrial and military applications, Acromag PCI and cPCI I/O boards deliver dependable operation in hostile environments.

COTS-grade components provide superior performance, and our extended temperature models will withstand -25 to 85°C ranges. Now *that's* reliability.

Reliability also means long-term availability. Acromag boards are designed for a life span of at least 10 years. Our lifecycle management program ensures that your system designs will be sustainable long into the future.

- Wide variety of long/short PCI, 3U/6U cPCI
- High-performance and high-density
- Software support for RTOS and Windows®

For flexible, capable, and reliable I/O solutions, think Acromag.

Acromag®
THE LEADER IN INDUSTRIAL I/O

www.acromag.com
800-881-0268 or 248-624-1541

Since the I/O interfaces appear as system master devices, maintaining coherency—ensuring data accessed by the external masters is valid—is an important consideration. Some communications processors require the system software to account for this, which can consume precious processor cycles during data flow processing. Communications processors with hardware-maintained cache coherency eliminate this potentially "buggy" software and improve performance by eliminating the need to run this software during data flow processing.

Some communications processors support preemptive cache allocation, also known as stash caching. Designers can take advantage of the fact that most data plane processing requires only a small portion of the protocol data unit (PDU), usually the beginning or header. The I/O unit is configured to force a predefined portion of the PDU into cache as it is received. This ensures that when the software starts to process the PDU, the required information is already in the cache, reducing the latency required for the processor to figure out how to process and convert the PDU. This feature can significantly increase data plane performance.

Another consideration is the use of packet processing resources. Some communications processors include one or more packet-processing cores. These cores typically handle MAC layer functionality, since a general-purpose processor cannot efficiently perform this task. The MAC functional interface within a communications processor is usually engaged by the general-purpose processor through easy-to-use, efficient ring buffer management of transmit and receive buffers for traffic flow and a configurable set of parameters for initialization. Minimal general-purpose processor overhead is required to transfer the traffic between the media interface and processor memory space. Some communications processors also have the option of performing forwarding without general-purpose processor intervention, giving the designer the flexibility to leverage both the packet processor and/or the general-purpose processor for forwarding traffic through the system.

High-performance communications processors may also possess specific data plane hardware acceleration features. These blocks target specific functionality and are typically used with simple parameter initialization and/or ring buffer management techniques, as in the MAC example. In contrast, network processors require specific programming or reprogramming to achieve similar functionality. Acceleration features include encryption/authentication, CRC validation and generation, traffic management/queuing support, content processing and classification co-processing.

Communications processors have a huge market lead over network processors because of the ease of programming and extensive knowledge base behind them, as well as the virtually unlimited flexibility associated with a communications processor-based solution. They have been used successfully for programmable data plane applications for several years and continue to evolve to meet the performance requirements of these applications. ◢

Freescale Semiconductor
Austin, Texas.
(512) 895-2000.
[www.freescale.com].

EMBEDDED PLANET™

design.
The next generation of connected devices.

develop.
Your products based on our platform.

deploy.
Your solution faster.

design. develop. deploy.

Are You Board With Me?

Start with an AMCC PowerPC 405 Processor.

Specify your I/O requirements.

Choose your operating system.

Relax while Embedded Planet delivers your rugged and reliable system.

4760 Richmond Rd / Cleveland, OH 44128
Tel: 216.245.4180 / Fax: 216.292.0561
www.embeddedplanet.com

EMBEDDEDPLANET™

• Quality and reliability by design • Long term supply •

PERFORMANCE IN THE EXTREME

Intel® Pentium® M Processor embedded SBCs for demanding environments

-30°C Power On

MAT 1100

- Intel® Pentium® M Processor 1.4, 1.8, 2.1G
- 5V only PICMG 1.0
- -30°C option
- Ideal upgrade for Pentium® III and Pentium® 4 designs
- Far lower power than the Pentium® 4
- Intel® Extreme 2 Graphics
- Optional High Performance Graphics using ATI 9000, dual screen, DVI/DFP, LVDS
- 2 Serial, 4 USB, dual Gigabit Ethernet, Sound, CompactFlash

12 to 36 watts - for fanless sealed system design

MAT 1111

- Intel® Pentium® M Processor 1.4G (738), 1.8G or 2.1G
- Intel® Celeron® M Processor 600Mhz
- Low power 12 watts to 36 watts depending on CPU and loading
- EBX size, Mini PCI, PC/104-Plus, CompactFlash
- Heat spreader plate to support sealed system design
- Fanless option
- Extended temperature -30°C power on option
- Intel® Extreme 2 Graphics, LVDS, CRT
- High Performance Graphics using ATI, dual screen, DVI/DFP
- Two independent video input channels; mux up to 4 inputs
- 5V only

Two Philips 7130 for PAL/NTSC input

Contact us

For full specifications go to our website

www.embedded-pc.com

Europe & rest of world:
Microbus plc
Tel: +44 (0) 1628 537333
Fax: +44 (0) 1628 537334
email: sales@microbus.com
www.embedded-pc.com

USA & Canada:
Microbus Inc
Tel: +1 (800) 688-4405 or (281) 568-4744
Fax: +1 (281) 568-4604
email: sales@microbus-usa.com
www.embedded-pc.com

Designed and manufactured in the UK

Microbus
Better by design

Microsoft

©2005 Microsoft Corporation. All rights reserved.

Microsoft, Windows CE and Windows Embedded are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

We Put the Power of Windows Embedded in our Robots

Each day, device developers like Robert Bonin are harnessing the power of Microsoft® Windows Embedded to build great devices, like the latest industrial robots from Chaveriat Robotique.

Microsoft® Windows CE provides hard real-time capability to their robots and powers a user-friendly handheld device which enables factory floor personnel to make program changes on the fly. With these and other innovations based on Windows CE, Chaveriat Robotique has reasserted itself as a serious force in the world of robotic arms, with robots that are more functional and less expensive.

What's more, Windows Embedded offers the timesaving tools, operating system technologies and thousands of drivers device developers need—so that they have the power to truly innovate.

> ***"We chose Windows CE because it offers real-time and graphics at the right price."***

— ROBERT BONIN / **Research & Development Manager**
Chaveriat Robotique / France

The Power to Build Great Devices—get it with **Windows CE, Windows XP Embedded,** or **Windows Embedded for Point of Service.**

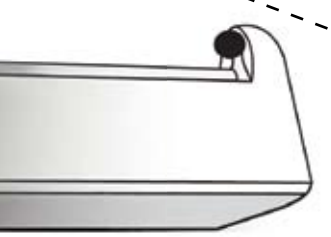

www.learnaboutembedded.com/robots2

Windows® Embedded

Beyond the Network Layer: NP-Based TCP Offloading via TCP Splicing

TCP performance of low-end devices in today's networks continues to be limited mainly by memory and CPU speeds. A TCP offload mechanism is proposed that takes advantage of TCP splicing and implements the offloaded functions in a remote network processor.

by Surekha Peri and Pravin Pathak
Agere Systems

The Transmission Control Protocol (TCP), a transport layer protocol, was originally designed for low-speed terrestrial links. However, with the proliferation of high-speed transmission mediums, as well as sophisticated access mechanisms, extensive computing power and memory are required for software-only TCP implementations.

In recent years, there has been an explosion in the speed of networks, CPUs and memories. Since the speed of Ethernet has increased much more rapidly than that of CPUs or memories, the performance of today's networks continues to be limited mainly by memory and CPU speeds.

One method for overcoming these limitations and accelerating a TCP/IP connection is TCP offloading. TCP processing complexity is offloaded from the host CPU to specialized TCP accelerators. Usually, the accelerating TCP offload engine (TOE) is a dedicated subsystem co-located with the host CPU. However, an alternative, more scaleable approach to achieving the same end, which also works with legacy equipment, is to move the complicated TCP/IP processing to a network processor (NP). This is a special-purpose programmable hardware device, which is connected to the TCP server over a high-speed link.

The NP-based TOE combines the low cost and flexibility of a general-purpose processor with the speed and scalability of custom silicon solutions. Furthermore, the NP offloads both memory-intensive and CPU-intensive processing from the TCP server. In the NP-based TOE mechanism described here, TCP splicing is used to enhance TCP performance.

An Alternative TCP Offload Approach

For low-end enterprise products limited by cost and size, addressing the memory/CPU bottleneck by introducing a remotely located TOE on a central networking element is a scaleable and cost-effective solution. This is especially true when con-

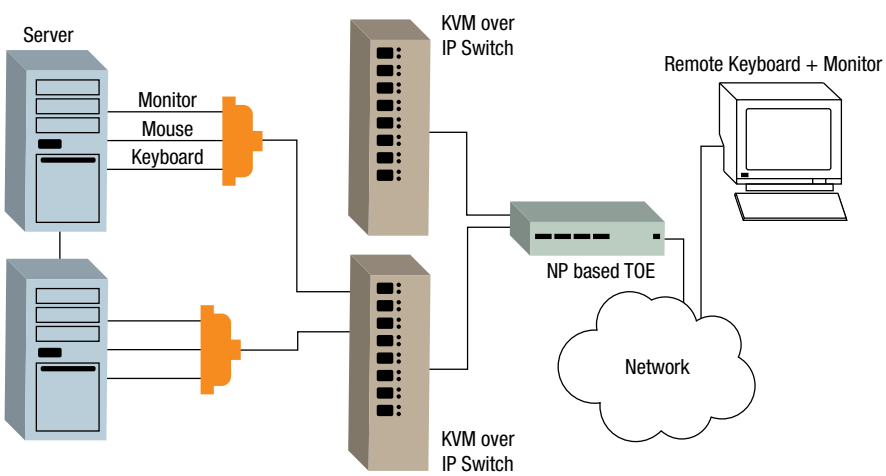

Figure 1 In a sample scenario, a network processor-based TCP offload engine substantially enhances the performance of low-end devices. Server performance is independent of the characteristics of the network between the network processor (NP) and the client and instead, reflects the characteristics of the link between the NP and the server.

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

trasted with the solution of increasing the capability of the host itself, such as via keyboard, video and mouse (KVM) switches.

In one sample deployment scenario (Figure 1), a remote keyboard, monitor and mouse control multiple, remotely managed servers connected to a KVM-over-IP switch. The mouse and keyboard events from remote controlling computers are transmitted over the Internet via the KVM switch to the server. Compressed monitor data is sent to the remote computer by the server via a KVM switch. The KVM switch offloads TCP processing to the NP-based TOE, simplifying KVM switch hardware.

TCP Offloading Using TCP Splicing

TCP offloading can be performed across two network entities, an enterprise TCP host and a remote NP. This is achieved by using a split TCP connection, also known as TCP splicing, a well-known technique for enhancing TCP performance. In TCP splicing, two independent TCP connections span a session: one from the client to the NP, and the second from the NP to the server.

Data from the server is locally acknowledged by the NP, thereby reducing server buffering requirements and speeding up congestion window growth (Figure 2). The NP buffers the data until the acknowledgment (ACK) from the far end (client) arrives. In the absence of an acknowledgment, the NP performs timer management to retransmit TCP segments toward the far end. The usage of local acknowledgments also shields the TCP server from any network congestion and excessive delays between the NP and the client. This reduces the memory burden at the server.

The NP also performs optional network address translation (NAT) for data from the server to the client involving checksum computations. Although this functionality is computationally intensive and heavily loads a general-purpose processor, an NP is well suited for such operations.

In contrast to traditional NP applications such as routing, the TCP offload application requires inter-packet dependency and connection-based state machines. Furthermore, this application deals with transport layer functionality, which is beyond traditional NP link- and network-layer processing.

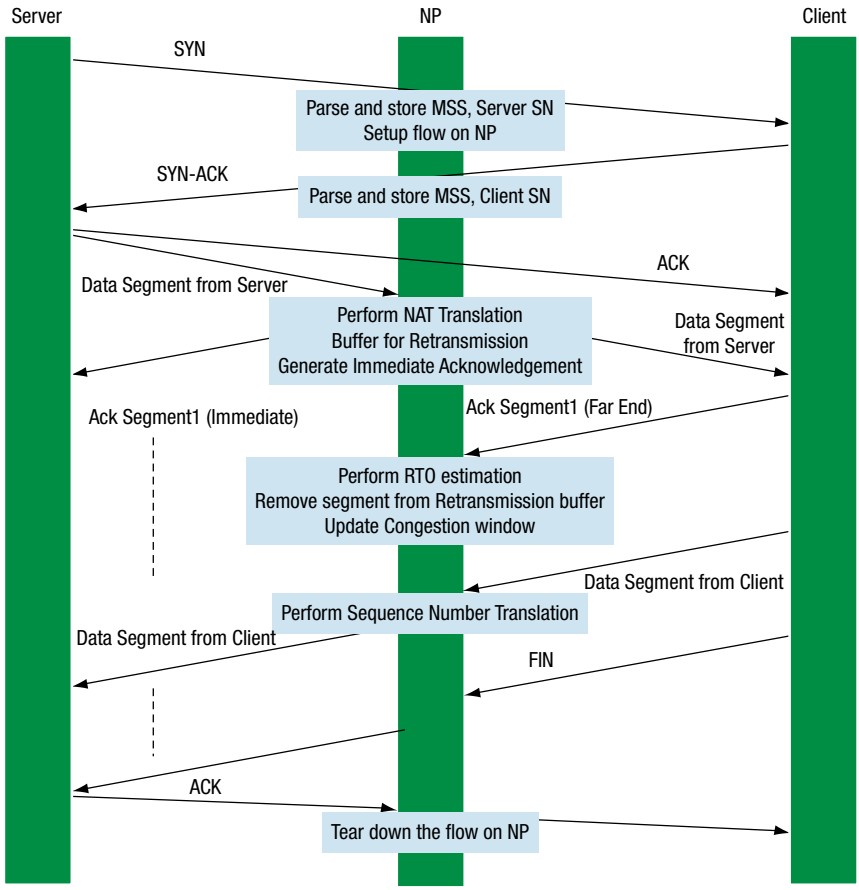

Figure 2 In TCP splicing, a well-known technique for enhancing TCP performance, two independent TCP connections span a session: one from the client to the NP, and the second from the NP to the server. TCP splicing on the NP-based TOE is shown here.

There are two fundamental complexities in TCP processing. Buffer management consists of storing unacknowledged segments, out-of-order segments and stalled segments in congested networks. Timer management involves Retransmission Timeout (RTO) estimation and triggering retransmission of unacknowledged segments. Both of these are offloaded to the NP.

The alternative TCP offload mechanism, using TCP splicing, is also useful when a high-performance server or network operates in tandem with a long-delay wireless and/or low-bandwidth network, as is predominant in defense and mobile networks. In this scenario, the TCP offload on the NP shortens the slow start phase as a result of the perceived shorter delay caused by the immediate acknowledgment. Note that the slow start phase dominates application performance when transaction sizes are small.

TCP congestion control is always tuned for congestion in wire line segments and performance degrades in error-prone wireless channels. TCP offload onto the NP prevents window collapse on the server, since error recovery is restricted to the NP and the client. This improves overall throughput.

Implementation of TOE with TCP Splicing on an NP

In order to achieve a proof of concept of this proposed alternative TCP offload mechanism, it was implemented on an Agere APP340 NP (Figure 3). The APP3xx series of NPs offers up to 2 Gbits/s throughput. The device contains a classifier; a traffic manager consisting of a scheduler, a buffer manager and a stream editor (SED); a state engine; and an on-chip control processor, i.e., an embedded host. Slow path, or control and management,

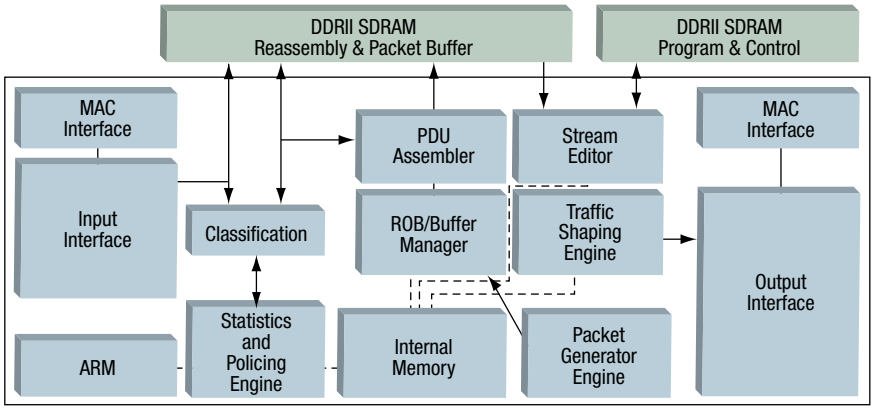

Figure 3 The proposed alternative TCP offload mechanism was implemented on an Agere APP340 network processor. The device contains a classifier, a traffic manager, a state engine and an on-chip control processor (embedded host).

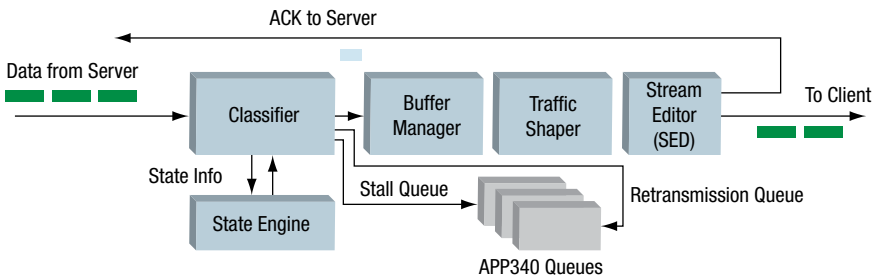

Figure 4 Processing of data segments from server to client on the NP-based TOE.

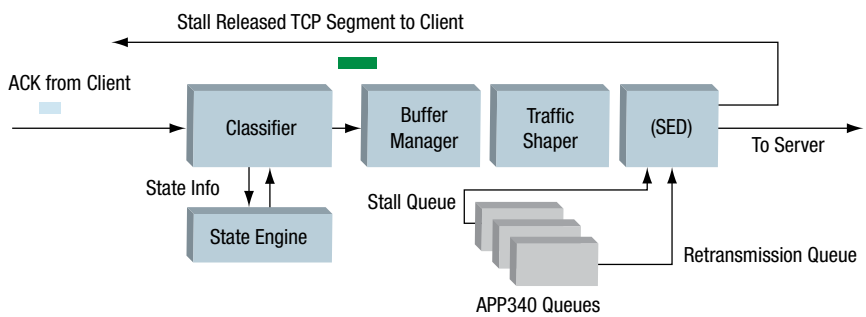

Figure 5 Processing of data from client to server on the NP-based TOE.

processing is achieved on the embedded host. Fast path, or data-plane, processing is performed by the rest of the components.

The classifier is used to identify the connection and its state. The state engine maintains the TCP state variables accessible to various functional blocks on the NP. The traffic manager runs the TCP congestion control protocol and is used to make scheduling decisions. The SED is used to perform sequence number manipulation and network address translation traversal.

The APP340 NP allows the traffic shaper functional block to control the scheduling of packets in each queue. For every flow, the NP maintains multiple destination queues, including primary transmission to the client, retransmission of packets to the client and transmission to the server.

The NP also supports hierarchical scheduling. This enables the use of auxiliary queues carrying the control information at the same level of hierarchy of each destination queue. These control packets are internally generated by the packet generator engine (PGE) based on the state machine. They are used to trigger scheduling mechanisms such as holding packets in the scheduling queues, and release them only when the TCP state machine requires them to.

The ability to control scheduling of TCP segments using this hierarchical scheduling architecture, along with the ability to maintain inter-packet state dependencies, makes the architecture of this NP well suited to TCP processing.

TCP Connection Establishment

When packets arrive at the NP's ports, the classifier block determines the flow of the packet, based on the TCP port numbers and the IP address. If the flow is not already present, the packet is forwarded to the embedded host for connection establishment.

The initial packets correspond to the TCP three-way handshake (SYN, SYN-ACK and ACK). Upon receiving these packets, the host parses the TCP options and obtains parameters such as sequence number (SN) and maximum segment size, and passes them to the state engine and the SED. It also updates lookup trees on the classifier that associate a flow identifier with the TCP connection.

The packets are then returned to the classifier for reinsertion to the packet stream. The outgoing TCP handshake packets are subject to NAT traversal and bypass the remaining TOE functionality outlined in the subsequent sections. Similarly, the host handles connection termination. The classifier parses the FIN field, which indicates termination of the connection, and forwards the packets to the host.

TCP Data Flow from the Server

When TCP segments for established flows enter the classifier, the flow identifier is obtained using the lookup tree. The APP340 NP is a block-based processor that employs a two-phase classification. In the first phase, individual blocks are processed. The re-assembled protocol data units are handled in the second phase. If a segment arrives out of order, it is held in the first pass re-order buffer in the classifier until the in-order segment arrives.

For in-order segments, TCP states are examined to see if the current sliding window size *(cwnd)*, the receiver's advertised window (rwnd), the memory restrictions of the NP and the size of packets in flight allow transmission of this packet. If the packet is to be transmitted, it is forwarded to the SED for NAT translation and transmission to the far-end client. A copy is stored in the retransmit buffer on the traffic shaper (TS). Another copy is sent to the SED for generating an immediate local acknowledgment. The SED uses incoming TCP/IP headers, the SN of the last ACK and the last received byte number to generate the local ACK.

If the TCP sliding window is stalled, the TS stops scheduling packets. The packets are effectively stored in the TS until the stall is released (Figure 4).

TCP Data Flow from the Client

When an ACK arrives at the classifier, the

classifier extracts the acknowledgment number and forwards it to the state engine. The state engine updates the *cwnd*, reflecting the slow start or the congestion-avoidance phase of the TCP flow. The engine also updates the *rwnd*, the size of unacknowledged data and the stall status of the sliding window.

A trigger is generated to the TS to remove acknowledged packets from the retransmission queue. The TS schedules all acknowledged packets and associated packets from the partner control queue, which are then discarded at the SED. If the TCP sliding window was previously stalled, a trigger is generated to the TS to resume servicing the TCP segments (Figure 5). This acknowledgment is then terminated at the NP.

If the client packet carries piggybacked data, the state machines are updated as described previously. The acknowledgment number in the header is modified to reflect the most recently generated local acknowledgment number and forwarded to the server. Furthermore, the next expected SN from the client is updated in the state engine to be used as the SN for subsequent local acknowledgments.

Retransmissions and Retransmission Timeout Estimation

For every packet in the retransmission buffer, a partner control queue generated by the internal PGE holds the transmit time of the segment. When an acknowledgment arrives, the round-trip time is computed as the difference between the time of arrival of the acknowledgment and the transmit time. The retransmission timeout (RTO) estimation is then performed as per IETF RFC 2988 recommendation.

When packets are placed in the retransmission queue, the PGE generates a control packet consisting of current time and retransmission time, or current time plus RTO.

The PGE is programmed to generate a periodic trigger. Each time this occurs, the TS examines the retransmission time. For every packet with an expired timer, the packet is scheduled to be delivered to the SED and the corresponding control packet is flushed. This is repeated until all packets with expired timers have been serviced. A copy of every retransmitted segment is looped back to be stored for further retransmissions. The fast retransmission algorithm as specified in the IETF RFC 2001 is achieved by the PGE generating a retransmission trigger upon receiving three duplicate acknowledgments from the client.

The NP-based TOE enables server performance to be independent of the characteristics of the network between the NP and the client. Instead, performance reflects the characteristics of the link between the NP and the server. In addition, memory requirements of the server remain independent of network congestion, and of the round-trip delays in the client network.

TCP performance of a low-end device can be substantially enhanced by using a TCP offload mechanism at a remote NP. Furthermore, the server is shielded from the resource requirement variability associated with various client environments. In addition, the congestion characteristics and link loss characteristics of the high-performance client network, along with long delays, are absorbed at the NP-based TOE, substantially enhancing the user experience.

Agere Systems
Allentown, PA.
(800) 372-2447.
[www.agere.com].

WHEN FAILURE IS NOT AN OPTION
Cool RoadRunner 4: PCI-104 Module with Intel® Pentium™ M745 processor
Toucan: COM Express Module with Intel® Core™ Duo processor L2400
COM Express
Embedded Systems Conference SILICON VALLEY
If your application needs reliable and ruggedized computers - call us!
LiPPERT Embedded Computers, Inc.
5555 Glenridge Connector, Suite 200
Atlanta, GA 30342
Phone (404) 459 2870 · Fax (404) 459 2871
ussales@lippert-at.com · www.lippert-at.com
LiPPERT
the embedded PC Company

Hardware Platform Management Building Blocks Speed ATCA and AMC Product Development

Developers of ATCA and AMC products often choose to use externally acquired platform management building blocks so they can focus their own scarce development resources on unique added values. Such building blocks must meet certain key requirements.

by Mark Overgaard
Pigeon Point Systems

Suppliers and consumers of telecom equipment are aggressively adopting the AdvancedTCA (ATCA) and AdvancedMC (AMC) architectures defined by PICMG. One key reason is that these architectures include an advanced hardware platform management infrastructure based on the Intelligent Platform Management Interface (IPMI). This infrastructure is a key enabler for the stringent service availability requirements that are typically mandated in telecom equipment and the mix and match interoperability that is crucial to the promised reductions in capital and operating expenses.

Implementing the relevant aspects of this platform management architecture in each type of ATCA/AMC component is not a trivial matter, however. Even the challenge of thoroughly understanding the management aspects of the relevant specifications is substantial. The IPMI specification is over 400 pages long; the management parts of the ATCA and AdvancedMC specifications account for hundreds of additional pages.

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

Building Block Candidates

Figure 1 shows the logical elements of an example ATCA shelf, focusing on the platform management aspects. An ATCA shelf manager communicates inside the shelf with IPM controllers, each of which is responsible for local management of one or more Field Replaceable Units (FRUs), such as boards, fan trays or power entry modules. Management communication within a shelf occurs primarily over the Intelligent Platform Management Bus (IPMB), which is implemented on a dual-redundant basis as IPMB-0.

ATCA boards that are AMC carriers add another level to the management hierarchy. Carrier IPMCs handle on-carrier management, including communication over a local IPMB-L with the Module Management Controllers (MMCs) that provide onboard management for the AMC modules.

An overall System Manager (typically external to the shelf) can coordinate the activities of multiple shelves. A System Manager communicates with each shelf manager over an Internet Protocol (IP)-capable transport, usually Ethernet.

As shown in Figure 1, the key interoperability interfaces in the ATCA/AMC architecture include IPMB-0, which connects the shelf manager (or more precisely its shelf management controller or ShMC) to the IPM controllers in the shelf. In addition, IPMB-L connects a Carrier IPMC to the MMCs on the AMCs installed on the carrier. There is also a system manager interface, by which a system manager communicates with one or more shelf managers.

The ATCA/AMC management architecture uses IPMI-based messaging across all these interfaces. PICMG has extended IPMI with dozens of commands to cover the unique needs of a modular platform designed for high availability. A key goal of this archi-

tecture is to ensure that independently implemented building blocks can communicate seamlessly across these interfaces, enabling coherent management of the entire shelf.

Compliant and Interoperable Across Key Interfaces

To achieve the necessary compatibilities, management building block suppliers must be thoroughly familiar with the specifications they're implementing. Ideally, they are directly involved in the PICMG subcommittees that create the specifications. Although there are some moves in that direction, there are currently no formal compliance processes for PICMG specifications. In the meantime, building block suppliers must thoroughly test their products for compliance.

In addition to internal compliance testing, independently implemented management building blocks must be compat-

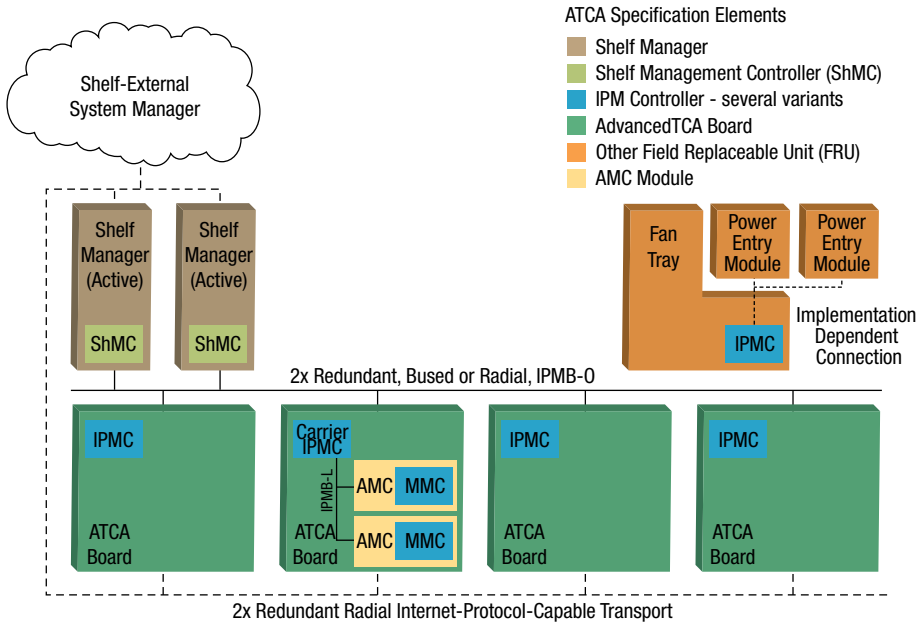

Figure 1 All types of management controllers in an ATCA/AMC-based shelf are building block candidates.

ATCA/AMC Interoperability Workshop Scenarios

AIW test scenarios developed by AIW participant companies systematically test interoperability of ATCA/AMC products across the key interoperability interfaces.

Interoperability Interface Focus			Management-oriented Test Scenario Categories (some categories with multiple test scenarios defined)
System Manager	IPMB-0	IPMB-L	
X			Remote Management Control Protocol (RMCP) interactions with shelf manager
X			Locating/validating/updating Shelf FRU Info
X			Shelf status and control
	X		Power on scenarios
	X		Life cycle of FRUs – states and transitions, payload and IPM controller control
	X		IPMB-0 redundancy, saturation, hang detection and recovery
X	X		Temperature event handling and fan control
	X		FRU addressing
	X		Point-to-Point and bused E-Keying for shelf fabric connections
	X		LED control
X			ShMC redundancy
X	X		IPMI entity handling
	X	X	AMC Sensor Data Record (SDR) repository maintenance
	X	X	AMC life cycle – states & transitions, Carrier IPMC and MMC resets
	X	X	AMC addressing
		X	IPMB-L operation
X	X	X	Message bridging to AMCs

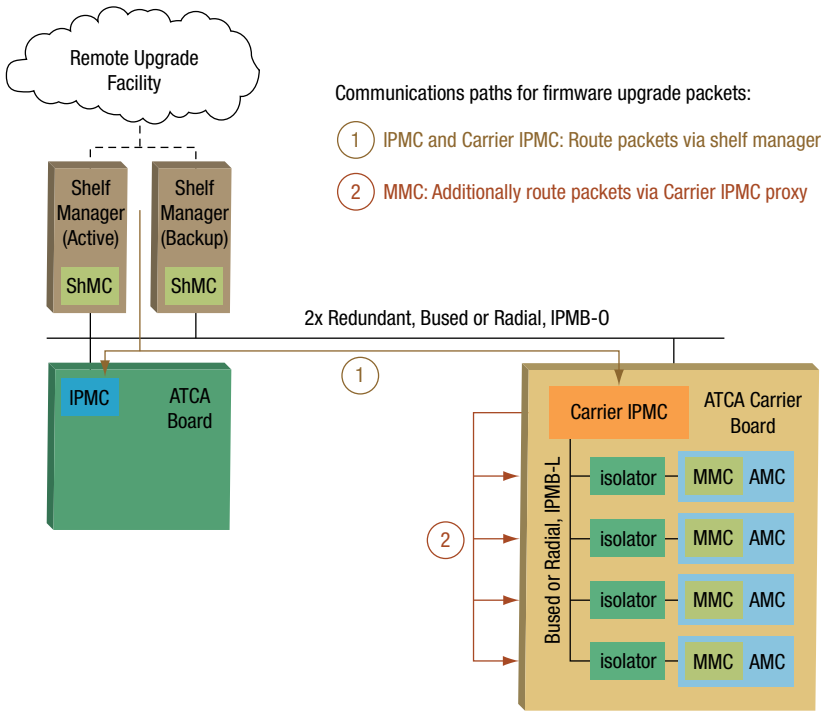

Figure 2 Firmware upgrade packets can be forwarded on IPMB-0 to IPM controllers and, for MMC upgrades, forwarded again on a carrier's IPMB-L.

ible across the interoperability interfaces shown in Figure 1. Since 2002 (even before the ATCA specification was initially adopted), PICMG has been running ATCA/AMC Interoperability Workshops (AIWs) to help ensure this compatibility. AIWs happen about three times a year and typically involve 15-30 PICMG member companies. Each company brings instances of their products for testing. During the AIW, participants systematically apply dozens of test scenarios (see sidebar "ATCA/AMC Interoperability Workshop Scenarios") with different combinations of components. For instance, each distinct board type is tested with each distinct shelf manager type.

Experience in AIWs has directly resulted in improvements in the ATCA and AMC specifications. For instance, independent implementers interpreted early specification requirements for FRU state transitions differently. Those and many other similar areas of potential incompatibility are now clarified in the specifications, but more will undoubtedly be found as testing gets broader and more intensive. Management building block suppliers must actively participate in AIWs to help ensure interoperability of their components and their customers' products as well as to help uncover any specification ambiguities that remain.

Adaptable Hardware and Software

Management building blocks must be adaptable in both the hardware and software dimensions. For instance, building blocks could be supplied with reference schematics for appropriate hardware components and source code for the corresponding software components. Of course, it is also important for the largest possible fraction of board developers to be able to use the reference design and firmware with only minor changes.

The article, "Accelerating the AdvancedTCA Ramp with Adaptable Shelf Management," in the May 2005 issue of *RTC* magazine covers this topic for shelf managers. For example, differing shelf requirements may mandate different placements of dedicated shelf manager FRUs in a shelf. A shelf manager building block needs to be adaptable in

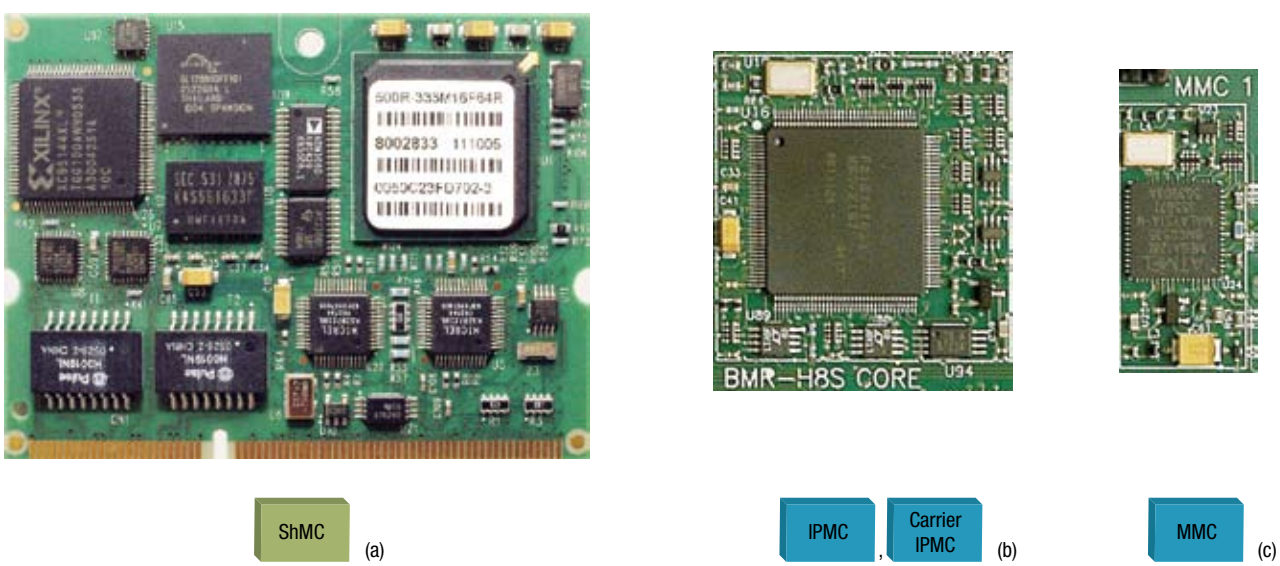

Figure 3 IPM Sentry management building blocks: (a) 67.6 x 50.8 mm SO-DIMM form-factor ShMM-500 provides the core of the ATCA shelf manager, ready for mating to shelf-specific carrier; (b) IPMC or carrier IPMC core, delivered as schematic and firmware source code (29 x 29 mm Renesas H8S variant pictured); (c) MMC core, delivered as schematic and firmware source code (24 x 12 mm Atmel AVR variant).

this and many other respects.

At the ATCA/AMC board/module level, adaptability is also critical, since the range of complexity and application focus for these boards is very wide. For instance, different developers and their customers may have quite different philosophies regarding the number and type of IPMI sensors (e.g., voltage or temperature) needed on a board. An IPMC building block must provide good flexibility in this area.

As another example, consider the E-Keying feature of the high-speed fabrics implemented in the ATCA/AMC architectures. Here, data structures stored as non-volatile FRU information associated with each board and shelf component indicate what protocols and widths of interfaces are supported. The shelf manager checks this data and issues commands to the communicating nodes that enable only compatible ports on each node. An IPM controller needs to communicate with the fabric devices on its board and enable or disable their ports, based on these directions from the shelf manager. How it does this communication may differ greatly between a node board that only attaches to two fabric ports and a multi-fabric switch board that attaches to dozens of fabric ports.

Reliable Firmware Upgrade

Management firmware intended for use in highly available ATCA systems must be reliably upgradeable without a physical visit to the equipment. Reliable upgrade typically implies that a separate copy of the upgraded firmware is downloaded to the IPM controller in such a way that the original firmware is preserved. If anything goes wrong at any time during the upgrade process, the original firmware can be restored and the IPM controller can go back to using it while the difficulties in the upgrade are sorted out.

The IPMI specification reserves a set of command codes for firmware upgrades, but doesn't define any specific command code points or associated semantics. Therefore, the firmware upgrade protocol by which commands using these codes implement an upgrade is currently specific to the developer of the IPM controller firmware. Work is underway within PICMG to standardize this process, but the results are not yet available.

IPMB-0 is the only mandatory interface to an IPM controller so it is a natural choice as a primary upgrade interface. Upgrade protocol packets can be forwarded from (and corresponding responses forwarded back to) an upgrade facility that either resides with or remotely accesses the shelf manager, which can play a proxy role by forwarding the packets over IPMB-0. Figure 2 shows this process. Figure 2 also shows how upgrade protocol packets can reach the MMC on an AMC, using the Carrier IPMC as a second-level proxy.

Management Building Blocks for ATCA/AMC

One set of ATCA/AMC management building blocks is the IPM Sentry family. Figure 3 shows some key family members and type(s) of in-shelf management controller targeted by each one. IPM Sentry management solutions are designed to address the requirements areas highlighted in earlier sections. IPM Sentry component developers are deeply familiar with the relevant PICMG and IPMI specifications; they are active in specification evolution, contributing substantial content.

At the shelf level, the SO-DIMM-sized ShMM-500, when mated with the shelf-specific carrier board, can be placed flexibly within a shelf; reference schematics for carrier board are provided. At the board level, reference schematics are integrated directly into the design of the managed board to optimize customization, cost and board real estate requirements.

Shelf manager source code is available, but there is also shelf-specific code for most shelves. Board management firmware source code is available to all board developers, but many can limit their changes to just a single source file to establish configuration details for their boards. At the shelf level, the ShMM-500 includes hardware support for reliable firmware upgrade. If newly loaded firmware does not confirm its operation quickly, hardware automatically reboots with previous firmware. At the board level, BMR reference designs provide an option to save existing firmware before new firmware is loaded. If new firmware doesn't work, automatic restore and reboot on previous firmware occurs. ◢

Pigeon Point Systems
Scotts Valley, CA.
(831) 438-1565.
[www.pigeonpoint.com].

Analogue & Micro Ltd

PowerPC Modules & Development Boards

USI offers a wide range of PowerQuicc™ development platforms and embedded modules from Analogue & Micro, available with Linux and eCos operating systems.

Modules and boards are available for the MPC850, MPC852T, MPC866T, MPC870, MPC875, MPC880, MPC885, MPC8250VR, MPC8270VR, MPC8280, MPC8247, MPC8248, MPC8271, MPC8272, MPC8343, MPC8347, MPC8349, MPC8541, MPC8555, & MPC5200.

Call us today!

USI Ultimate Solutions, Inc.

10 Clever Lane
Tewksbury, MA 01876 USA
Toll Free: 866.455.3383
Fax: 978.926.3091
Email: info@ultsol.com
Web: www.ultsol.com

ATCA Shelf Management for High Availability

The ATCA architecture provides centralized hardware management for high-availability systems. ATCA Management Controllers are specifically designed to work within the ATCA architecture.

by Kevin Hart
Vadatech

The increasingly complex systems in use today require increasingly sophisticated management tools. Even a desktop computer benefits from power and fan management, and the benefits increase for larger systems. The larger and more heterogeneous a system becomes, the more data must be gathered and organized to properly maintain it. Intelligent Platform Management Infrastructure (IPMI) is one standard that has been developed to address these issues.

The central player in the IPMI architecture is the management controller (MC). An MC, typically a microcontroller, monitors sensors and provides access to identifying information for one or more components. An MC runs independently of the component that it monitors. A centralized baseboard management controller (BMC) communicates directly with system management software, and allows the operator to access all of the management information in the system. The BMC also provides a system event log (SEL) that stores critical events received from other MCs.

IPMI defines a set of messages that can be sent from one MC to another. Many of the messages relay sensor or configuration information. Other messages are specific to communication between the BMC and system software. Still other messages are used to relay message data from one chassis to another. IPMI messages may be sent using a number of protocols, of which the most important is the Intelligent Platform Management Bus or IPMB. IPMB messages are sent over an I^2C bus.

Using IPMI, an operator can securely establish a connection to a local or remote BMC. He can examine serial number and version information for installed components, and get a list of recent temperature or voltage spikes from the SEL. He can configure the BMC to page him if a failure condition or other event of interest occurs. The system software to accomplish this is the same, regardless of whether the system is a desktop computer or a multichassis data center.

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

ATCA Enhances IPMI

IPMI is a broad specification, but there are areas it does not address. It has only minimal support for power management, and does not support hot-swap devices at all. It is primarily a protocol and data specification, and so is not useful for hardware integration. The AdvancedTCA (ATCA) specification from the PCI Industrial Computer Manufacturers Group (PICMG) was designed to address these issues.

The ATCA specification focuses on modular computing components. It defines chassis, shelves and node boards. A chassis holds one or more shelves. Each shelf has a fan tray, a redundant pair of power supplies and slots for up to 16 node boards. Node boards may provide power and signals to rear transition modules. The ATCA specification defines connectors, voltage levels and physical dimensions for all of these components, so that components from different manufacturers will interoperate. It also extends IPMI by adding messages to support hot-swap, power control and electronic keying, or E-Keying.

ATCA defines a shelf management controller, or ShMC, which provides much the same functionality as a BMC in an IPMI system. It controls and monitors all the subsidiary MCs in a shelf, and consolidates the configuration information

and sensors for all of the managed components. It communicates with the other MCs via dual-redundant IPMB, which can be either a single bus or a radial configuration (Figure 1). System software communicates primarily with ShMCs. ATCA also defines an Intelligent Platform Management Controller, or IPMC.

An IPMC typically resides on an ATCA node board, which it controls and monitors. An IPMC relays sensor events to the ShMC, and performs power and E-Keying functions in response to messages from the ShMC. A related PICMG document, the Advanced Mezzanine Card Base Specification, describes a type of node board called an AMC Carrier, which supports up to eight hot-swappable AMC modules.

The IPMI specification allows third parties to define additional types of sensors as needed. ATCA defines a hot-swap sensor type to facilitate hot swapping. Every hot-swappable entity has a hot-swap sensor associated with it. This sensor generates an event when the entity transitions from one hot-swap state to another. For example, when an ATCA node board is inserted into a running shelf, the IPMC on the node board sends a hot-swap event to the ShMC. The event message indicates that the ATCA node board has transitioned from the "not present" state to the "present" state. Although it is normally implemented in software, the hot-swap sensor can be read just like any other IPMI sensor.

Electronic Keying

ATCA defines a very flexible system for interboard I/O. The majority of the pins that connect an ATCA node board to a shelf are differential signal pairs. The signals are connected to a backplane in the shelf and are grouped into channels that in turn are grouped into interfaces. There are three interfaces: base, fabric and update channel. The base interface is a 10/100/1000BASE-T dual-star network among two hub boards and the remaining 14 ATCA node boards. All shelves must implement the base interface, so node boards have a standard way to communicate with the outside world.

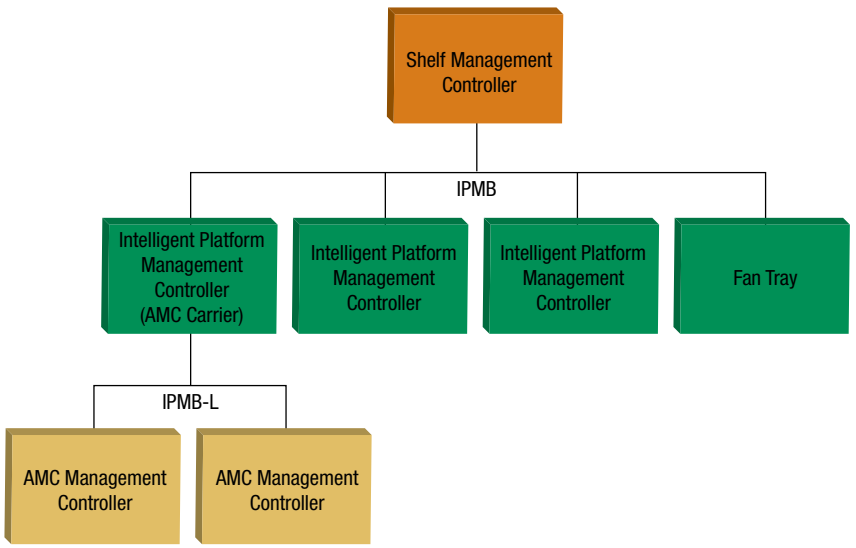

Figure 1 The shelf management controller communicates over the IPMB with management controllers on node boards. Management controllers are also implemented on AMC cards and communicate via the IPMC on the carrier card to the shelf manager via a local version of the management bus, IPMB-L.

The update channel interface consists of ten differential pairs, and is typically used to connect two adjacent node boards. It may be used for redundancy interlock or to allow two boards to exchange large volumes of traffic with no impact to the rest of the system. The fabric interface is a general-purpose interconnect, which can be routed in many different ways. The ATCA specification describes channel routing assignments for a variety of backplane configurations.

The channels in the fabric interface can potentially implement many different interconnect technologies. The ATCA specification lists Ethernet, InfiniBand, StarFabric and PCI Express, and refers to these as "link types." In addition, it provides a way for OEMs to specify new link types using globally unique identifiers, or

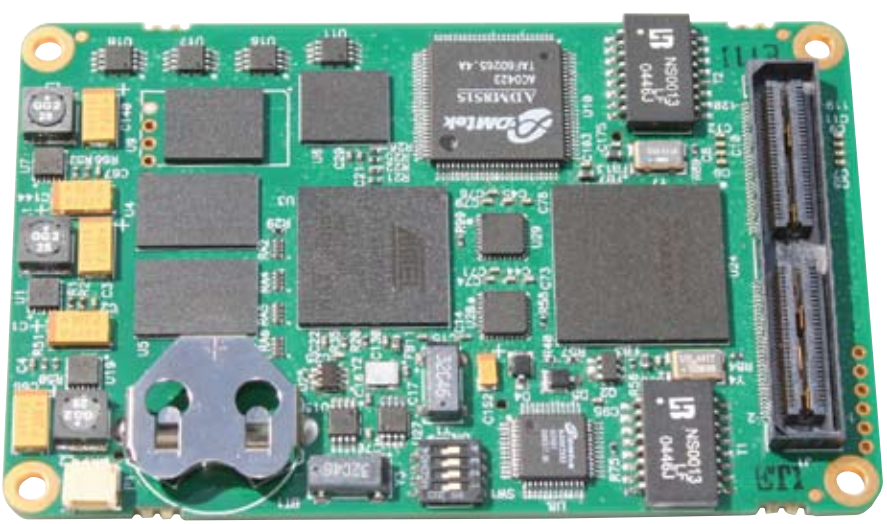

Figure 2 The Vadatech VT001 offers complete IPMI 2.0 hardware/software for shelf management, and blades in the size of a credit card and with a power dissipation of less than one watt.

GUIDS. A link type in combination with a set of channels is called a link. The ability to implement arbitrary links on a node board leads to potential communication problems between node boards that are using the fabric interface. The solution is electronic keying.

When a node board is introduced into an ATCA shelf, the ShMC queries the node board to determine what links it uses. The ShMC then uses its knowledge of the shelf's backplane configuration to determine which links are connected to compatible links on another node card. The ShMC sends commands to both node boards to enable every link that is connected to a compatible link. All other links on the newly inserted node board are explicitly disabled by commands from the ShMC. This scheme allows new interconnect technologies to be deployed safely without modifying the shelf hardware or software. A similar technique is used to interconnect AMCs on AMC Carriers.

As you can see, the ShMC has a lot to keep track of. It must manage state, power and sensor information for up to 16 node boards and over a hundred AMCs. It must deal with E-Keying and log and forward events. It must adjust fan speeds in the fan tray according to temperatures reported by node boards. It must be able to deal with one or more encrypted streams of messages from system software, and manage the associated user and password information. Node boards require varying levels of complexity from their management controller, with AMC Carriers on the most complex end of the spectrum. AMC management controllers (MMCs) have the lowest level of complexity, though this depends somewhat on the AMC payload.

Supporting ATCA

Would-be manufacturers of ATCA hardware face a familiar decision. On the one hand there are the integration issues and licensing fees associated with using a third-party management controller. On the other hand are the cost and risk of implementing a management controller from scratch. Either approach can divert resources from the main development effort. The relative importance of the costs will vary from product to product. Licensing fees for example are more important in a low-cost, high-volume product.

Vadatech's management controllers, for example, range from the VT001 for shelf or VT020 for AMC Carrier management (Figure 2), to the VT027 and VT026, for node board and AMC Module management, respectively. These management controllers can be configured and integrated by the manufacturer or by Vadatech. The VT001/VT020 software was developed in C++, and was designed to accommodate different hardware configurations without requiring a change to the executable binary. The configuration files provide a layer of abstraction from the hardware. Internally, extensive modularization and abstraction of sensors and other devices provides another layer. The configuration files assign a unique set of resources to each process.

Vadatech
Henderson, NV.
(702) 896-3337.
[www.vadatech.com].

And you thought SHE was sensitive?

What about your embedded system?
Just like sensing the pea under a mattress, your A/D and D/A system must detect and measure accurately too. So, start your system design with boards that handle real world irritants like temperature, drift and electrical noise.

Sleep easier knowing your measurements are accurate because the effective number of bits for the whole system is close to the ideal 12, 14, 16, or 24-bit resolution of the A/D silicon. Benefit from onboard fault protection circuits that ease system integration by handling the noise and power sequencing challenges common to multi-board designs. Eliminate the need for cumbersome and complicated layers of auto-calibration software with our low drift specs designed to keep your measurements constant over time.

And our wide selection of cost sensitive A/D and D/A boards is something that will make any princess happy.

We invite you to check out our 12, 14, 16 and 24-bit A/D and D/A single board computers from 12MHz to 533MHz, including VGA, CompactFlash,® Ethernet, GPS, CAN, wireless, and the operating system of your choice.

Call now and let our technical team help smooth out those sensitive issues in your system.

MICRO/SYS

3730 Park Place, Montrose, CA 91020
Voice (818) 244-4600 Fax (818) 244-4246
info@embeddedsys.com

www.embeddedsys.com

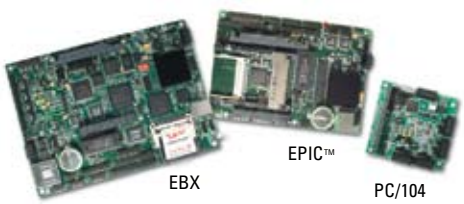

24-bit A/D
Our MPC624 provides 24-bit A/D for PC/104 with four channels of instrument grade voltage measurement and/or four channels of direct connect to off-the-shelf sensors.

16-bit A/D, 14-bit D/A
Our headless Pentium,® the SBC2596, offers Ethernet, GPS, CAN, DC/DC converter, and CompactFlash on EBX form factor with full 16-bit A/D and 14-bit D/A.

14-bit A/D, 14-bit D/A
Our SBC4495 is a 486/586 EPIC form factor with 14-bit A/D and D/A, VGA/Flat Panel, PCMCIA, wireless, and GPS.

For more fully-featured A/D and D/A SBC boards, visit our web site.

©2004 Micro/sys. All brand and/or product names listed are registered trademarks of their respective owners.

Windows Embedded Windows Server System

OEM Technical Seminars

presented by Microsoft

Everyday developers are choosing Microsoft® Windows® Embedded for their next great device.

See for yourself how Microsoft Windows Embedded can help you optimize device design, enhance reliability and integration, and improve time to market and development costs.

These sessions include a presentation and technical demonstration of the following Microsoft Windows Embedded Products, plus time to network at the exhibition.

- Windows CE 5.0
- Windows XP Embedded with Service Pack 2
- Windows Embedded for Point of Service
- Windows Embedded Server Products
- Special Guest Presentation by intel

Win a FREE Portable Media Center!

register today www.microsoftembedded.com

Attend this event and receive a FREE evaluation copy of Windows XP Embedded or Windows CE.

2006 Strategic Locations

Mountain View, CA January 24	Reston, VA April 11
Redmond, WA February 24	Boston, MA May 23
Chicago, IL March 21	Atlanta June 20

Exhibitors

Windows Embedded Partners

ADVANTECH Applieddata.net Embedded Computer System Ardence ICOP LOGIC

Microsoft Authorized Embedded Distributors

Arrow Electronics, Inc. AVNET Bell Micro bSQUARE

*Content, topics and exhibitors may vary by city. Please visit www.microsoftembedded.com for complete information.

an RTC Group production

Executive Interview

RTC Interviews Mads Lillelund, CEO, Carlo Gavazzi

***RTC*: Now that we've turned the calendar page over to 2006, and you've been in the saddle as president of Carlo Gavazzi Computing Solutions for coming close on a year now, what is your overall impression of the market and where do you see the embedded computer business going?**

Lillelund: I am encouraged about the market and the growth prospects for our businesses. Within the embedded computer business we see growth opportunities across our four targeted segments, which are: Military, Telecom, Industrial and Medical. Within the military segment we see continued growth with VME-based products as well as a keen interest in learning about new technologies. We have also seen this segment be an early adopter of innovative cooling technologies. Within our Telecom segment, the key driver is the move to ATCA and now an early interest in MicroTCA.

Clearly with new technologies coming into the market and more innovation and value required within the chassis, customers are increasingly requiring companies like us to provide increased design and engineering talent, quality and to help them fast track products to market. These requirements play to our strengths, so we are very excited about our prospects.

***RTC*: Carlo Gavazzi Computing Solutions comprises several different groups that have been consolidated under the Computing Solutions banner. Can you briefly tell our readers how the divisions evolved and what areas of the technology and expertise they cover?**

Lillelund: Carlo Gavazzi Computing Solutions has grown through a combination of organic growth as well as acquisitions. We will continue this strategy as we bring new technologies to the market. I am a big believer in a strategy that combines technology alliances, partnerships and key acquisitions. By executing well on this strategy we can fast track products to the market and bring new technologies to our customers in an effective and cost-efficient manner.

Today we have three key groups that are uniquely positioned in the market, but that also allow the ability to leverage each of the groups' strengths with each other. The three groups are: Electronic Packaging & Embedded Computing (formerly known as Mupac), Fabric & Connectivity (formerly known as Aurora Technologies) and Channel Access.

Electronic Packaging & Embedded Computing is focused on the continuous participation in the development and application of the PICMG, VITA and PIC SIG Open Architectures. Basically we apply this open architecture knowledge in designing system-level packaging solutions for our customer applications in the Telecom, Military, Industrial and Medical segments.

Fabric & Connectivity is focused on off-the-shelf solutions based upon StarFabric and PCI Express to address the needs of PCI and PCIe system expansion. The traditional Aurora brand serial products are still being maintained along with our ControlTower console management system. The rapid adoption of PCIe on motherboards and the reduced number of full height and/or length PCI slots have expanded the need and available market for reliable, well-designed, scalable system expansion. Our solutions allow people to maintain their existing investment in PCI cards, drivers and software and control their migration path to PCIx and PCIe. I have also chartered this group to stay abreast of the emerging and maturing fabrics outside of StarFabric and PCIe including InfiniBand and RapidIO, while continuing their committee activities with the PCI SIG groups.

Channel Access is a technical sales channel targeted toward the enterprise space. Our skill sets center around our technical sales ability and best-in-class product offerings, which feature NMS Communications, PIKA Technologies and our Aurora Technologies (Fabric & Connectivity) products. We are aiming to grow this business by extending the product lines and offering integrated systems for the enterprise space.

Our key driver in everything we do is to constantly innovate for our customers. Fortune 1000 companies as well as start-ups and medium-sized companies look to us to provide key capabilities to them. Chief among them are design and engineering, quality products, speed to market and in many cases they also take advantage of our integration capabilities.

***RTC*: ATCA has been very much on the front burner for a number of companies yet it still eludes the kind of market penetration that many analysts have fore-**

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

cast. Will 2006 be the year that ATCA achieves momentum and starts to reach the volumes that will make it a leading bus architecture? Why or why not?

Lillelund: As I prefer to spend my time with customers, I have closely followed their increased interest to the point that we now see ATCA starting to build momentum. We are now seeing traction in the market and orders from our customers in the telecom/datacom segment. We have therefore positioned the company to provide our customers with the best-in-class ATCA portfolio on the market. To further drive innovation, we just announced a partnership with Comtel, where we will combine our technical experience to provide further innovation for our customers.

***RTC*: In a similar vein, the military is becoming increasingly network-centric and as such, is requiring more network-oriented hardware. We've heard of at least a handful of ATCA systems and subsystems already going to military prime contractors and for development of military systems. As a supplier of hardware to the military, have you witnessed an increase in activity of ATCA in the military? If so, for what applications and what kind of equipment?**

Lillelund: Yes, the military seems very keen on understanding the impact of ATCA on their current and future applications. We have in many instances arranged for informational sessions on ATCA with our military customers. In fact, one of the first custom backplanes that we did was for a contractor that serves the military segment. Clearly there is an attraction to ATCA due to the high bandwidth that ATCA provides. High bandwidth is critical to future applications within the military such as live field simulations in 3D. The processing of unencrypted data within a single ATCA chassis is another application that is being explored in the military space.

***RTC*: MicroTCA—the system platform based on AMC mezzanine cards—has been criticized from a number of sources for a number of reasons such as connector cost and robustness, inject and eject mechanism, limitations to board thickness and overall fragile construction. Reports now, however, indicate that it has passed shock and vibration testing. Do you think MicroTCA will emerge as a winning solution for a variety of systems (Lucent already envisions systems with hundreds of AMC cards) or will it default to some other standard? Which one(s) do you see in the running?**

Lillelund: There is going to be great demand for AMCs because of the flexibility they provide to system designers. How these AMCs will be deployed (in MicroTCA or ATCA) comes down to where the common system cost break occurs between the two platforms.

Even though the MicroTCA standards are not finalized, we have several customers who want to move ahead with early MicroTCA developments. We therefore launched a MicroTCA Starter Kit at the Bus & Board show, which we will provide to early adopters.

***RTC*: VME has been a staple in Carlo Gavazzi's electronic packaging stable lineup for many years. However, the business is changing radically. What was the number one VME supplier has gone end-of-life on many of its designs and according to reports has backed off on most new development. What was the number two maker of VME hardware has abandoned the architecture and was acquired by what was the number one player—the combined company now focusing on ATCA and communications strategies. While the remaining VME market remains relatively strong, 1) can it continue to be a viable force in the embedded computer board and subsystem market in light of growing competition and 2) will a critical mass of suppliers continue to offer the variety of products required? 3) Is there a market for VME outside the military?**

Lillelund: First, with the amount of VME design-ins, legacy products and VME design knowledge that is available in the market, VME will continue to be a viable force in the embedded computer board and subsystem market. Secondly, yes, for the reason stated in my pervious answer. And finally, yes, we still see many VME solutions being used in Industrial applications.

***RTC*: Switched fabric technology appears to be something that's going to happen—it's just a matter of when. Several variants from ATCA and AMC, which we just discussed, above to other approaches such as VITA 36 and VITA 41 are beginning to emerge as is CompactPCI Express. Do you see any of these emerging as critical technologies over the next several years? Which one(s)? How will the emergence of hybrid and non-VME products impact the backplane and packaging business?**

Lillelund: Let's face it. Switched fabric technologies have happened, although we are just at the beginning. We have several CompactPCI 2.16 solutions shipping at production levels and many ATCA systems moving from development to deployment. We have seen some interest in VITA 41 as well. The move from traditional bus-based architectures to fabric-based architectures has extended the life of the board and backplane, or should I say blade and backplane method of electronic packaging for at least a decade. The next technology challenge is cooling.

***RTC*: The Carlo Gavazzi Fabric & Connectivity division was an early supporter and developer of StarFabric. And while we believe StarFabric offers a lot of advantages compared with some of the other switched fabrics on the market, it never achieved a significant market presence, and as a standard does not have a critical mass of vendors. Does the StarFabric family of products you make continue to make sense in a world dominated by PCI Express at the low end? Does Carlo Gavazzi plan to continue its family of StarFabric boards or do you have plans for a new product family based on another fabric?**

Lillelund: Our early adoption of StarFabric and fabric interconnects in general continues to reap benefits. The risk we took did not have the direct financial return we expected, but the experience has been invaluable. PCI Express is a derivative of StarFabric in many ways, and with the substantial adoption that PCIe has already achieved, we are well positioned

to provide PCIe-based products. One key application of StarFabric that we continue to have success with and will continue to support is PCI expansion. We have a 1U PCI expansion system that uses StarFabric as the enabling technology. This product has generated more interest recently due to the rapid changes in available PCI slots at the motherboard and system levels. At the board level, we will continue support of both a PCI and PMC bridge and will introduce a low-profile version this quarter. These cards are used as host interfaces for the PCI expansion systems and have had limited success in general applications. We will not continue support of our StarFabric/2.17 centralized switch due to lack of demand and success of 2.17.

***RTC*: Back in 2000 and 2001, Carlo Gavazzi was on the acquisition trail buying such companies as Aurora and Channel Access among others. However, since then, it has slowed down. Do you believe the entire embedded computer industry has slowed its acquisition rate? Why? High valuations? Too few companies available?**

Lillelund: Some of the slowdown in mergers and acquisitions within the industry was probably due to the burst of the bubble and most companies regrouping at the time. However, you have seen some consolidation in the market recently. I think this trend will continue in 2006 and 2007. The companies that have come through the downturn in market, but have focused on strengthening their financials, will be the ones that are looking for acquisitions to aggressively grow their businesses. Carlo Gavazzi Computing Solutions has a strong balance sheet and cash position. We intend to be looking at assets that either bring in new technology platforms or increase our market share in key segments.

***RTC*: The general theme of this year's Bus & Board conference was applications. And one of the application areas where we've been seeing a great deal of activity is in the convergence of voice, data and video, both wireless and wireline (fiber optic) on the same network. And, on the cable side of the business, Time Warner, Comcast, Cox and others are pushing in the same direction. Joining the fray are wireless companies Sprint Nextel and Internet companies Google and Yahoo. Cisco has already begun deploying some hardware, but it would seem a natural for ATCA, MicroTCA and other standard form-factor systems. What application areas have you seen the most activity in recently? Have you noticed any increase in companies supplying equipment for converged communications services? From the ATCA packaging side? From the fulfillment division side?**

Lillelund: The packaging business is definitely seeing increased traction in the areas of VoIP, Internet Media Servers and other network-intensive applications. These systems are being deployed on CompactPCI 2.16 and ATCA. uTCA will also play in the space when it gets finalized. Increasingly, bandwidth-intensive applications will strain the available technologies. We will continue our rich history of innovating to remove the roadblocks from our customers so that they can focus on their core competencies.

Shatter
Processing Limits in Your
Wireless Applications
Breakout performance from BittWare's New
DSP Advanced Mezzanine Card (AdvancedMC™)
Combining the best-in-class performance of Analog Devices TigerSHARC with the benefits of AdvancedMC, the B2-AMC is ready for your wireless applications: 2G, 2.5G, 3G, WiMAX, and Software Defined Radio. Using the superscalar architecture of the ADSP-TS201 TigerSHARC DSPs, the B2-AMC provides an unprecedented 14.4 GFLOPs and 57.5 GOPS of processing power. A large Altera Stratix II FPGA implements BittWare's ATLANTiS architecture, seamlessly integrating the DSP processing power with your switch fabric of choice: Serial RapidIO™, PCI Express™, Advanced Switching Interconnect, GigE, or XAUI™ (10 GigE). In addition to the multiprocessor scalability of the TigerSHARC, the board features a variety of front and back panel I/O interfaces, an IPMI system management interface, and system synchronization via AMC system clocks. Ideal for use in AdvancedTCA®, MicroTCA, or custom systems, the B2-AMC will give your wireless application the power it needs.
Go ahead...
Shatter your limits.
BittWare
www.bittware.com
TigerSHARC DSP is a registered trademark of Analog Devices, Inc.

Data Recording SoC Aims at Space-Constrained Apps

Using separate boards for data acquisition and recording is cumbersome and expensive, especially in applications where weight, power and size are critical. A new system-on-chip (SoC) developed by TEK Microsystems and QinetiQ makes both functions available on a single board.

Configured as advanced firmware in an FPGA, the JazzStore SoC is available as an option on five different Tekmicro/QinetiQ RTES data acquisition or processing products. These include the Neptune dual-channel 2 Gsample/s data acquisition board, the Venus five-channel 105 Msample/s data acquisition board, the Triton 2 Gsample/s data acquisition with 2 Gsample/s digital waveform generation board, the Callisto FPGA processor VXS board (5 Virtex ii Pro 50) with 12-channel high-speed digital I/O ports, and the JazzFiber four-channel fiber optic PMC module.

Intelligent operations to trigger and record block sizes of 8-32K samples and other typical data recording features are implemented with a resolution greater than 1 sample clock period. Xilinx Rocket IO on-chip high-speed serial interfaces connect the FPGA to a Fibre Channel disk through a media transceiver. Recording speeds can be greater than 64 Mbytes/s sustained via a Fibre Channel port. The JazzStore SoC is priced at $5,000 as an add-on option.

TEK Microsystems, Chelmsford, MA. (978) 244-9200. [www.tekmicro.com]

Rugged VME Blade Server Is IBM Clone

A new 6U VME blade server is a ruggedized version of the IBM PowerPC JS20 blade design. The PowerNode5 from Thales Computers provides high performance and full binary compatibility with the dual-PPC970 JS20.

The PowerNode5 processor SBC has two PPC970FX G5 CPUs operating at 1.6 GHz, with an operating temperature range of 0° to +55°C. It features up to 2 Gbytes of DDR SDRAM ECC memory and peak memory bandwidth of 6.4 Gbytes/s as well as peak PCI memory bandwidth of 850 Mbytes/s. For connectivity and expansion, the PowerNode5 includes a 64-bit PMC site, a USB port and two 10/100/1000BASE-T Ethernet ports. It supports the PPC64 version of Linux and includes an onboard Serial RapidIO switch fabric.

A pre-integrated version, the PowerMP5, is also available, with full data transport and management software. Both Red Hat Linux and Wind River VxWorks are supported. The PowerNode 5 is available in both conduction-cooled and convection-cooled versions. Pricing for the PowerNode 5 starts at $11,260.

Thales Computers, Raleigh, NC. (919) 231-8000. [www.thalescomputers.com].

Storage Blades Provide High Density

High-availability communications systems are demanding ever more storage capacity, especially as service providers continue to offer more value-add services. With this in mind, Performance Technologies has introduced two 6U storage blades that each offer 1 Terabyte of storage for embedded database, data caching and file serving applications.

The CPC5900 Storage Blade and CPC5910 Storage Expansion Blade are CompactPCI 2.16- and 2.9-compliant. The 6U CPC5900 is a high-density, high-availability NAS or SAN storage blade with two board-level and drive-level hot-swappable SATA hard drives. The onboard, 800 MHz PowerPC processor operates on Performance Technologies' NexusWare Linux development environment. The CPC5900 supports full RAID arrays (0, 1, 0+1, 4, 5, 6), and can be used as a PXE boot server for automated deployment of raw boards. Dual Gigabit Ethernet ports provide network connectivity, and 16 Mbytes of flash memory provide expansion space.

The CPC5910 features two board-level and drive-level hot-swappable, enterprise-class SATA hard drives with a SATA interface. When paired with the CPC5900, the CPC5910 supports both NAS and SAN applications. The CPC5900 is priced at $2,600 in volume, and the CPC5910 at $1,350.

Performance Technologies, Rochester, NY. (585) 256-0200. [www.pt.com].

IPv6 VMEbus Managed GbE Switch Platform Targets Mil Apps

A key requirement for network-centric warfare is the support of IPv6. A family of 6U VMEbus Gigabit Ethernet switches from GE Fanuc provides integrated, fully managed Layer 2/3 switching and also supports IPv6.

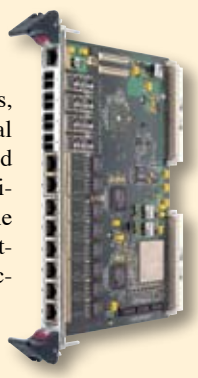

The new RM921 family is available with 12 (single slot) or 24 (dual slot) front-panel GbE ports, copper and/or fiber connectivity options, conformal coating and an optional 10 GbE uplink port. Based on the Motorola XPC8240 processor and the multi-layered Broadcom BCM56302 switch fabric chip, the RM921 is IPv6-compliant with full support for routing and tunneling, and also supports Layer 4-7 functionality. The switches are compatible with current IPv4 deployments and can be fully managed via SNMP, Telnet or a Web server interface.

Advanced QoS functionality includes per-port priority queuing, traffic prioritization and dynamic multicast filtering (IEEE 802.1p), and packet filtering. IPv6 features include IPv6 addressing and specification (RFC 2460), neighbor discovery for IPv6 (RFC 2461), stateless address auto configuration (RFC 2462), ICMPv6 (RFC 2463) and IPv6 over Ethernet (RFC 2464). Pricing starts at $4,000.

GE Fanuc Embedded Systems, Huntsville, AL. (800) 322-3616. [www.gefanuc.com].

Zero Latency Circuit Switch for VXS

Engineers building signal processing systems such as radar, sonar, SIGINT, electronic intelligence and real-time imaging sometimes may require only a few VXS boards interconnected with point-to-point connections in a custom VXS backplane using a simple protocol. However, this solution still demands the time and expense needed to develop a custom backplane.

The Phoenix CSW1 circuit switch from VMetro solves these problems by providing the ability to use off-the-shelf VXS backplanes and change the CSW1's routing topology. VXS payload cards can be connected together with fast, protocol-agnostic data links. These links can be arranged in any topology.

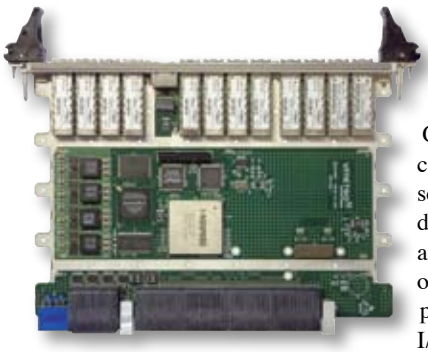

The Phoenix CSW1 supports 56 1x VXS backplane data links, each operating at up to 3.2 GHz, enabling intrasystem communications and remote sensor interfaces. Backplane data links can be utilized as 4x payload slot links in off-the-shelf VXS star backplanes. Additional optional I/O capabilities include up to 12 front-panel transceiver channels and a short XMC site. The Phoenix CSW1 is available 8 to 12 weeks ARO. Pricing begins at $7,995. Discounts may apply to customers using other Phoenix products.

VMetro, Houston, TX. (281) 584-0728. [www.vmetro.com].

Conduction-Cooled Audio Card Sports Front and Rear I/O

A conduction-cooled PMC audio card features a single PMC slot with front and rear I/O, widening the options for audio connections. The ccPMC481 from AcQ InduCom is based on the CrystalClear PCI audio interface controller and the CrystalClear CS4299A SoundFusion audio codec '97.

The ccPMC481 includes four audio connections via the front panel using 3.5 mm stereo connectors. Audio connections can be made via the rear I/O using the optional PIM481-J PMC I/O module, which mounts on a rear transition module and has four 3.5 mm stereo connectors. The audio interface consists of a line input, a microphone input, a line output and a speaker output with an onboard stereo amplifier.

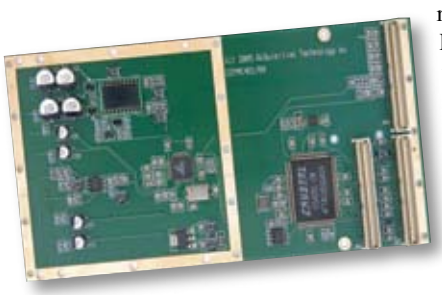

The card also includes a 33 MHz/32-bit PCI interface, supports 3.3V and 5V signaling and has a 5V and 12V power supply. Price is $640 each.

AcQ InduCom, The Netherlands. +31 (0)412 641922. [www.acq.nl].

VXS Extender Board Provides Easy Access

To provide full access for testing or debugging to a circuit card under test, extender boards must bring the card completely out of its card cage or enclosure. Until recently, VXS extender boards were not available because of the lack of a right angle receptacle for VXS. Elma Bustronic has developed its VXS Extender Board for hub slots by engineering a rigid-flex-rigid PCB.

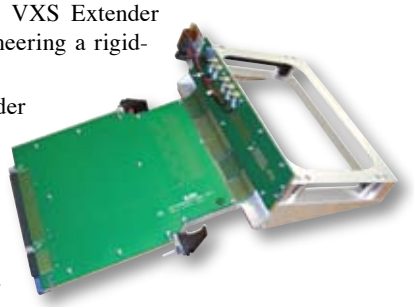

The Elma VXS Extender Board for hub slots utilizes a right-angle pin connector that plugs into the backplane, connected to a flex circuit wired to the straight receptacle to receive the plug-in board. There are test points for all of the differential pairs on the multi-Gigabit fabric connector for the hub slots. The VXS Extender boards come in a 10-layer stripline design for the rigid PCB and a microstrip design for the flex circuit portion. Also featured is an ampmeter, which measures the current and has a digital indicator on the front panel showing the status.

Lead time is 2-4 weeks ARO. Pricing is under $3,000, and depends on configuration and volume.

Elma Bustronic, Fremont, CA. (510) 490-7388. [www.elmabustronic.com].

PMC SATA Controller Manages Four Attached Devices

Although not yet mainstream in the embedded space, Serial ATA is gradually displacing older drive technology. A SATA controller from Technobox manages four attached devices.

The Versatile SATA Controller, P/N 4382, provides four 1.5 Gbit/s links on a PMC front panel. Additional features, supported by a PLD, are provided for rear I/O. The PCB mounted connector is a "latching style" version permitting use with SATA cables that have a built-in, positive-retention latch. The board employs an Intel 31244 controller, which connects the PCI-X bus to four SATA links. PCI-X can operate up to 133 MHz under ideal conditions, but the board will also operate in lower speed 33 MHz and 66 MHz buses and in 32-bit and 64-bit modes. Using the PLD, additional functions can be enabled, including an interface to an I²C controller, four UARTs, 5V-tolerant digital I/O, analog inputs and system monitor. A boot EEPROM automatically loads the PLD on power-up. A front I/O-only variant of the controller, P/N 4383, does not incorporate the PLD.

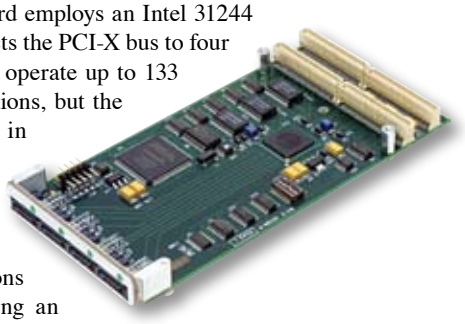

The fully populated 4382 lists for $495. A depopulated variant, the 4383, lists for $295. Quantity discounts are available.

Technobox, Mount Laurel, NJ. (609) 267-8988. [www.technobox.com].

Get Connected with companies and products featured in this section.
www.rtcmagazine.com/getconnected

Ruggedized TigerSHARC-Based 6U VME/VXS Board Sports Flexible I/O

Designed for demanding multiprocessor-based operations, and targeted at a broad range of applications including radar, sonar, communications and imaging, the T2V6 from BittWare is a 6U VME/VXS board featuring eight 600 MHz ADSP-TS201S TigerSHARC DSPs. It also incorporates two FPGAs, a PMC+ site, 512 Mbytes of SDRAM and flexible, high-throughput I/O interfaces via the front panel, VXS switch fabric and 2eSST VME. The T2V6 is available in both air-cooled and conduction-cooled versions and gives VME users full access to the power of the TigerSHARC. To take full advantage of the TigerSHARC processors, this multiprocessor board implements a scalable architecture, dubbed ATLANTiS, which enables designers to easily reconfigure the I/O and communications. Implemented in the FPGA, ATLANTiS can also be used to implement pre-, post- and co-processing. By tightly integrating the DSPs, VXS switch fabric, PMC interface and I/O peripherals with the onboard FPGA, ATLANTiS gives designers extremely flexible options for configuring, routing and processing high-throughput I/O.

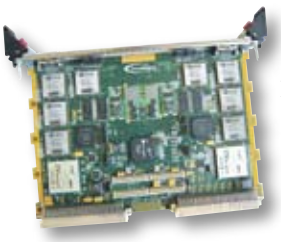

The T2V6 includes a VME-PCI Bridge supporting 2eSST, a PMC site with BittWare's PMC+ I/O extensions, high-performance digital I/O headers and sixteen SerDes high-speed serial transceivers (eight to the VXS switch fabric, eight to the front panel). The commercial (air-cooled) version is priced at $9,995 in OEM quantities. A fully rugged conduction-cooled/conformal coating version will be priced at $11,995 in Q2 2006.

BittWare, Concord, NH. (603) 226-0404. [www.bittware.com].

Network Attached Storage Blade Saves Space and Reduces Single Points of Failure

A hot-swappable 6U, single-slot removable network attached storage (NAS) module is targeted at a variety of embedded applications requiring highly available and redundant blade-level storage. The RAIDStor model A-8027 from ACT/Technico is suited for embedded file servers, boot servers and Voice-over-IP, as well as NAS embedded storage using a network file system (NFS). In a system using multiple single board computers (SBCs), each with its own drive, the boot device would be a single point of failure.

The new NAS RAIDStor can be used with either PICMG 2.16 or VITA 31.1-compatible systems, since these standards use dual star network topology. The blade is managed as an appliance via a Web browser interface or via SNMP interface. RAIDStor provides two packet switched backplane-compatible (PSB) Ethernet connections to implement links to two different networks. When used in a PSB-compatible backplane, an Ethernet connection is made to two fabric switches, allowing redundant links to connect other node boards within the chassis. Other features include automatic network synchronization of data written to an active RAIDStor operating in a mirrored NAS configuration, as well as support of redundant NAS controllers, allowing data to be mirrored between RAIDStors over the network. Pricing starts at $6,295 for a standard RAIDStor blade.

ACT/Technico, Warminster, PA. (215) 956-1200. [www.acttechnico.com].

Market leading Intel XScale® based single board computers...
Combining the very best of low power Intel XScale RISC technology and the industry standard PC/104 format...
VIPER
• 400MHz PXA255 Intel® processor
• Ultra-low power 1.9W (typical)
• TFT/STN graphics controller
• 64Mbytes DRAM / 32Mbytes Flash
• Ethernet, USB and serial ports
• CompactFlash (CF+) and PC/104 bus
• Operating temperature range:
- -20°C to +70°C (standard)
- -40°C to +85°C (industrial)
• Ready-to-run development kits for:
- Windows® CE 5.0 - Embedded Linux
- VxWorks® - QNX Neutrino
VULCAN
• 533MHz IXP425 Intel communications processor
• x2 100baseTx Ethernet ports
• Accelerated encryption hardware
• x4 USB 2.0 and x4 serial ports
• 64Mbytes DRAM / 32Mbytes Flash
• CompactFlash (CF+) and PC/104 bus
• Operating temperature range:
- -20°C to +70°C (standard)
- -40°C to +85°C (industrial)
• Ready-to-run development kit for:
- Embedded Linux
CUSTOM DESIGN & BUILD
www.arcom.com
Pb
RoHS Pb
Arcom
us-sales@arcom.com
888-941-2224
Think Embedded. Think Arcom.

Get Connected with companies and products featured in this section.
www.rtcmagazine.com/getconnected

New Quad PowerPC SBC Debuts on VITA 46 (VPX)

A VITA 46 (VPX)-based single board computer, the 6U VPX6-185 from Curtiss-Wright Controls Embedded Computing, delivers a nominal backplane bandwidth of 8 Gbytes/s via four Advanced Switching Interconnect (ASI) ports and two PCI Express VITA 42 XMC/PMC sites. In addition, it offers an extensive list of standard features such as Gigabit Ethernet (GbE), serial ports, MIL-STD 1553 and mass storage interface options.

The VPX6-185's processing engine is the Freescale 8641 single/dual-core PowerPC device with dual integrated 64-bit memory controllers. For numeric-intense processing, the 8641 offers the Altivec instruction set extension, which performs up to 8 floating point operations per cycle. The dual XMC/PMC sites on the VPX6-185 conform to the VITA 42.3 standard and can accept either standard PMC modules or XMC modules featuring PCI Express. The XMC I/O connectors are provided with dedicated backplane I/O pins.

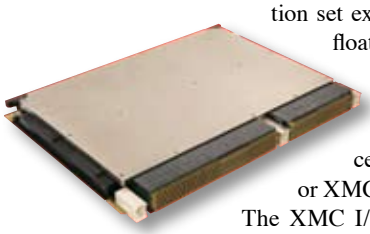

The VPX format provides backplane connectors capable of handling signaling speeds of up to 6.25 Gbits/s. The board's VPX connectors support the 2.5 Gbits/s signaling rate of the ASI ports and also allow multigigabit/sec I/O to be handled from the XMC/PMC sites. Software support for the VPX6-185 includes BSPs and driver suites for VxWorks 6.x with Wind River Workbench 2.0, INTEGRITY 6.x and Linux 2.6. Pricing for the VPX6-185 starts at $10,500. Evaluation units will be available in Q4 2006.

Curtiss-Wright Controls Embedded Computing, Leesburg, VA. (703) 779-7800. [www.cwcembedded.com].

6U VME Board Delivers Rugged Data Acquisition

A new data acquisition card that combines rear I/O with onboard signal conditioning circuitry in a rugged, single-slot package, the ICS-8145 from Radstone is built around the rugged 6U VME form-factor. Designed for harsh environment, confined space high-frequency sonar, wet end of towed arrays, torpedoes, autonomous underwater vehicle (AUV) and unmanned underwater vehicle (UUV) applications, the ICS-8145 is a complete, self-contained solution that eliminates the need for auxiliary equipment that characterizes other offerings.

The ICS-8145 features 2.5 Msample/s oversampling ADC, gain, anti-alias filtering and clock generation resources on the card, eliminating the usual need for auxiliary signal conditioning equipment. Because the card is expected to be deployed in space-constrained areas, all I/O is routed through the backplane. The differential analog inputs are routed though VME P0, while configuration and control is carried out via the VME 64 interface. An FPDP II interface (on the VME P2 connector) allows data from multi-card systems to be easily organized and collected. The inclusion of up to 8 Mbytes of onboard storage means that fast transient captures can be read back over the VME without data loss. Single unit pricing starts at $7,650.

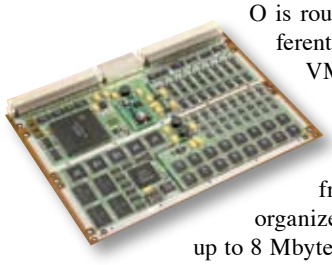

Radstone, Ottawa, Ontario. (613) 749-9241. [www.radstone.com].

Multi-Service AdvancedMC for High-Availability Telecom Edge Access Apps

Targeted at convergent gateway and ATM to IP interworking applications like convergent 3G wireless, Voice-over-Internet-Protocol (VoIP), media gateways, DSLAMs and switch/routers, the Telum 1204-O3 from SBS Technologies offers flexible functionality and performance in an AdvancedMC form-factor. The Telum 1204-O3 can independently interconnect between any supplied protocols, which can be selected on a per-port basis. For increased adaptability, a full set of IP services can be offered over any Layer 2 protocol including ATM AAL5, PPP and Ethernet. An intelligent AdvancedMC card, this module enables high-performance packet processing, directly terminating and switching traffic on the card rather than by sending traffic to an off-board processor. This enables system designers to build high-performance communications systems using a distributed processing architecture.

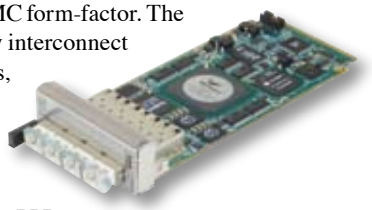

The Telum 1204-O3 is a multi-service card supporting ATM and IP interworking with the capability of processing up to 750,000 packets per second. This module supports local port-to-port Automatic Protection Switching (APS) to enable continuous service in the event an OC-3 link should fail. By using Small Form-factor Pluggable (SFP) connectors, the Telum 1204-O3 can operate in OC-3 multi-mode or OC-3 single mode for configuration flexibility. In keeping with the European Union's environmental policies, a RoHS version will be made available. Pricing starts at $2,799, and it is available now.

SBS Technologies, Albuquerque, NM. 505) 875-0600. [www.sbs.com].

RoHS-Compliant PC/104-Plus SBC Meets EU Requirements

The PC/104-Plus "Lynx" CPU board from VersaLogic meets the requirements of the European Union Restriction of Hazardous Substances (RoHS) directive (2002/95/EC). The board was developed after considerable research and testing with lead-free components, solder and new manufacturing processes. In addition, VersaLogic worked closely with its vendors to evaluate and certify the plating of all components in order to address tin whisker concerns in long-term applications.

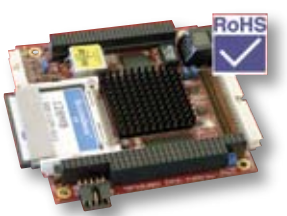

This board is targeted at vehicular/aircraft controls, medical equipment and other OEM applications in the European market. Since it does not include onboard video circuitry, it is most applicable to headless applications in process monitoring, data acquisition, communications routing or safety monitoring of critical systems. The Lynx features the AMD Élan SC520 processor CPU. Standard onboard features include 64 Mbytes soldered-on SDRAM, 4 COM ports, Ethernet, IDE, LPT, two counter/timers, floppy interface and keyboard/mouse ports. Standard pass-through connectors allow the board to be used either above or below PC/104 add-on modules. It may also be used as a CPU module for a larger system, by plugging it into a proprietary baseboard that includes user I/O circuitry. The Lynx, part number VL-EPM-4g, is currently available from stock. Pricing is $425 in OEM quantities.

VersaLogic, Eugene, OR (541) 485-8575. [www.VersaLogic.com].

Transceiver Development Platform Handles Wideband Recording and Playback

A fully programmable development platform that targets transceiver applications such as military radios and commercial wireless base stations, the RTX 2502 from Pentek incorporates playback capability, using the same data-capture front end as the company's RTS 2501 and adds four D/A channels with two up-converters for playback. The unit comes preconfigured as a data recording system with a graphical user interface (GUI) to make the RTS 2502 ready to use out of the box.

The combination of data recording and playback capabilities in a single unit means that instead of using simulated signals and signal analysis tools during system development of radar or communication systems, engineers can take the RTS system into the field to capture and generate real-world signals in real time for direct validation of both signal-processing algorithms and system hardware. Multiple Virtex-II FPGAs provide signal processing, formatting and other data manipulation for the input and output channels. Pentek supports developers in using the FPGAs with its GateFlow FPGA Design Kit and also offers GateFlow IP Cores that include FFTs, wideband and narrowband digital down-converters and a radar pulse compressor. An RTS 2502 base system starts at $29,995 for hardware only and $41,495 with bundled SystemFlow and GateFlow software.

Pentek, Upper Saddle River, NJ. (201) 458-2311. [www.pentek.com].

iSCSI Storage Blade Features SAS Technology for Data/Telephony Apps

A new, fault-tolerant, standards-based storage blade is designed for demanding applications requiring storage solutions that leverage the features of both iSCSI and Serial Attached SCSI (SAS) technologies on a single slot. The HighWire SAS iBlade from SBE addresses storage needs in the telecommunications market. SBE has ported the PyX iSCSI stack onto its HW400c/2 core processing platform and is bundling it with Astek's PMC-based SAS host bus adapter to create a complete, fully tested storage transport solution ready to plug into any CompactPCI networking and communications server or system.

The base platform of the SAS iBlade is a CompactPCI I/O processor featuring a PowerPC, up to 1 Gbyte SDRAM, two PTMC sites, dual Gigabit packet switched backplane (PICMG 2.16), front panel Gigabit Ethernet and H.110. The PyX iSCSI software ensures seamless data migration by employing a multi-channel data transmission approach. A PMC-based SAS daughter board from Astek delivers up to eight 3 Gbit/s SAS/SATA ports, integrated RAID (0, 1) and point-to-point connections for increased aggregate bandwidth, higher availability and enhanced reliability. The blade supports up to two SAS daughter boards to provide a total of 16 ports. The blade is available with support for Linux OS. Pricing starts at $6,000 and varies based on configuration requirements and OEM quantities.

SBE, San Ramon, CA. (925) 355-2000. [www.sbe.com].

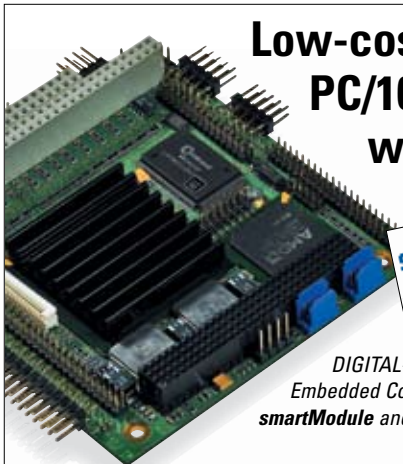

Low-cost, fanless PC/104 CPU board with 500MHz

Embedded Systems Conference Silicon Valley Booth 1538

*DIGITAL-LOGIC offers a wide range of Embedded Computer in **PC/104, EPIC, EBX, smartModule** and other form factors*

MICROSPACE® MSM800SEV

- AMD Geode™ LX800@0.9W Processor with 500MHz
- up to 1 GByte DDR-DRAM SODIMM
- 2x USB V2.0, 2x COM, LPT1, IrDA, FD, E-IDE, LPT1, KB, MS
- LAN Ethernet (boot from LAN)
- AC97 Sound
- integrated 2D graphic controller CRT up to 1920x1440x32bpp@85Hz 24Bit Digital Video Interface (TFT)
- 0°C up to +60°C (option: -40°C up to +70C°)

Advanced Digital Logic Inc.
4411 Morena Blvd. Suite 101
San Diego, CA 92117-4328
Phone +1 858 490 0590
Fax +1 858 490 0599
sales@adlogic-pc104.com
www.adlogic-pc104.com

Further information: www.digitallogic.com/msm800sev_en

DIGITAL-LOGIC
smart embedded computers

Get Connected with companies and products featured in this section. www.rtcmagazine.com/getconnected

RTC delivered straight to your inbox

Switch to Digital Delivery

RTC magazine now comes straight to your Inbox! Just visit RTCmagazine.com to convert your print magazine subscription to Digital Delivery. Each month your downloadable issue will include everything in the print magazine in a searchable PDF format. Archive your issues for portability and easy reference!

Check out all of *RTC*'s Digital Services

RTC Online Edition

Visit RTCmagazine.com and browse the latest issue of *RTC* magazine online. You'll find all the keen editorial insight, detailed technical articles and product info that make the print edition a must-read for industry insiders.

RTC Update Email Newsletter

Get the latest headlines and product news from RTC Update. The newsletter arrives in your Inbox mid-month, keeping you up to date and on top of our fast moving industry. Sign up at RTCmagazine.com today!

RTC Get Connected!

RTC's Get Connected will hook you up with the right people without a lot of hassle. When you need important product details right from the source, look for the Get Connected symbol in *RTC*'s print and digital versions.

Get Connected

RTC Supplemental Content

Looking for more info? *RTC* Supplemental Content has it. Select *RTC* magazine articles include additional content posted online. Dig a little deeper into the technologies shaping the future of our industry.

RTCmagazine.com

RTC

the RTC Group

Successful Carrier Card Design for COM Express

With all the really difficult high-speed CPU bus signal design done for them, carrier card designers can focus on the potentially new serial bus signal integrity design requirements, which differ from the parallel buses that they are more used to dealing with.

by Andrew Pitt
RadiSys

COM Express is a modular embedded computing standard developed for applications requiring higher CPU performance or higher I/O bandwidth, along with design flexibility in an extremely compact form. The standard defines mechanical, electrical and thermal requirements for a highly integrated COM mezzanine and specifies a rich set of high-speed serial I/O interfaces. This combination preserves current parallel interface technologies and ensures a smooth migration path for future upgrades.

The standardization and modularity offered with COM Express enables broader design reuse, reduced development expense and significantly less risk. This means faster time-to-market and time-to-revenue when developing embedded applications such as medical, test and measurement, gaming, entertainment, industrial automation and point-of-sale products.

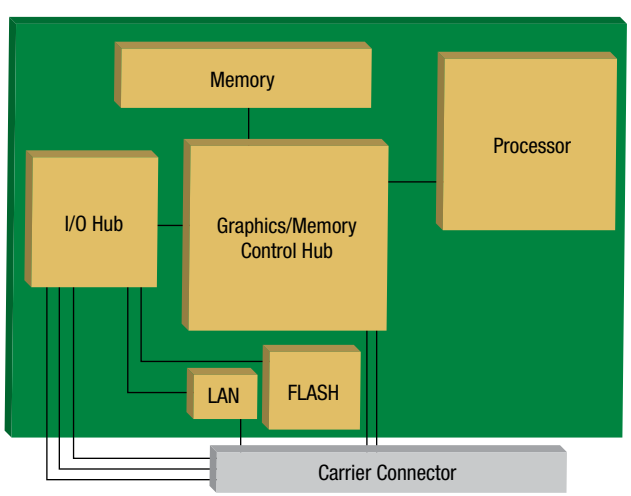

Figure 1 Typical COM Express module.

Get Connected with companies mentioned in this article.
www.rtcmagazine.com/getconnected

Developers have access to a variety of computer system modules that can be used to design application-specific carrier cards of any size and shape. Considerations for each are similar and include the physical and mechanical aspects, especially component height. In addition, careful attention must be given to electrical parameters such as impedance matching, termination, noise and transmission line loss. Then there are the issues of interfacing the carrier card to module control.

Mechanical Placement and Component Height

There are two carrier card stack options: 5 mm and 8 mm, which require different connectors (Figure 2). In a 5 mm stack, clearance between the carrier card and the module's PCB is 5 mm (DIM A). The standard specifies that components to be mounted on the backside of the module must have a maximum height of 3.8 mm (DIM B). Components placed on the carrier card topside are limited to a maximum height of 1 mm (DIM C) with the exception of mating connectors. Carrier card topside components of up to 1 mm allow a gap of 0.2 mm between the carrier card module and bottom-side components. This may not be sufficient space in situations where flex or severe vibration are expected.

With the 8 mm stack carrier card option, topside components are limited to 4 mm in height (DIM C) with the exception of mating connectors. Using carrier card topside components up to 4 mm again allows a gap of 0.2 mm between carrier card topside components and module bottom side components. A larger gap may be created by using a lower component height envelope than the one specified in the standard.

Electrical Considerations

Regardless of the bus, basic electrical considerations for carrier card developers include impedance matching and ground referencing for the differential pairs, signal termination and noise susceptibility. For impedance matching, all signal references are to ground, and traces should run over continuous ground planes. Differential pair traces susceptible to noise should be routed away from noise generating devices. Transmission pairs usually require termination if not used. However, USB signals require no termination. To ensure that trace lengths on the carrier card, and associated losses, are kept small, calculate the sum of the transmission line losses caused by modules, cards, connectors and cables.

Several examples are based on RadiSys carrier card design and are for reference only. Correct trace width and spacing for ideal impedance can only be calculated after the PCB stackup is known.

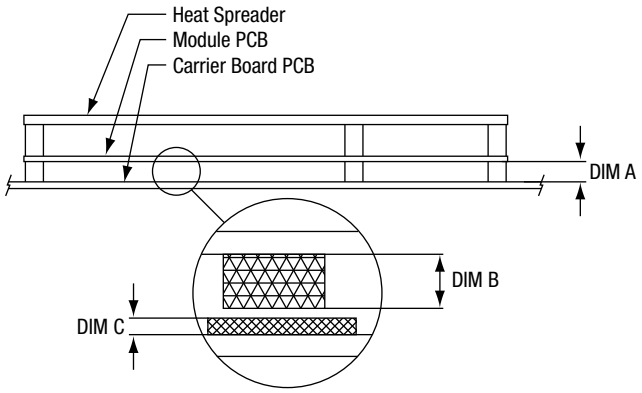

Figure 2 If parts are mounted on the top (DIM C), component height is a key concern in the design.

PCI Express and SATA

The following terms define the logical and physical configurations of the PCI Express bus providing required signal paths, of appropriate bandwidth, between devices.

- **Lane**: Four pins on the connector can be used for a single PCI Express transmit pair and a single receive pair.
- **Link**: A group of PCI Express lanes between two PCI Express agents. Allowable link widths are x1, x2, x4, x8, x16 and x32.
- **Bucket**: A group of PCI Express lanes on the connector. The 32 PCI Express lanes are divided into buckets describing how available PCI Express lanes should be assigned to connector pins.

The module allows configurations up to a maximum of 32 PCI Express lanes, giving the equivalent bandwidth of 4 x 10G Ethernet connections. Chipsets on COM Express modules have a variety of lane and link capabilities. On some, the PCI Express lanes can be grouped into various links under software control, while on others the links are a fixed width. Mapping the chipset lanes to the COM Express lanes and the grouping of the lanes into links is referred to as link configuration.

There are four main design issues that should be considered for carrier cards that use the PCI Express bus: All signals must be referenced to ground and all traces routed over continuous ground planes. Transmit-and-receive lanes should be routed as differential pairs and the differential impedance kept within the range of 100 ohms ±20%. When changing layers ensure trace matching for either transmit or receive pair, in the same layer. A grounded via should be placed within 250 mils of the layer changed point. Avoid routing traces under noise-generating devices because noise affects signals.

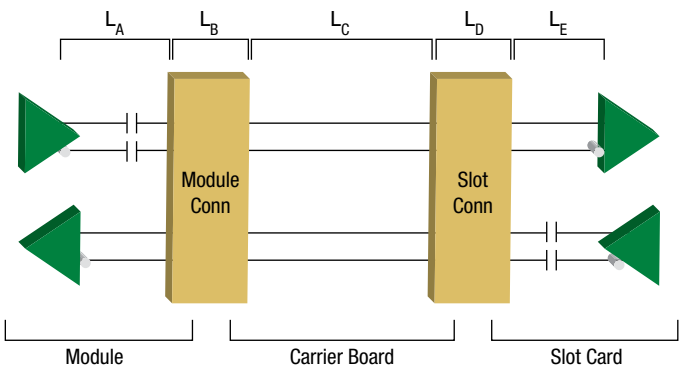

Figure 3 PCI Express insertion loss budget includes the module, carrier card, slot card and connectors (sum of = $L_A+L_B+L_C+L_D+L_E$).

Figure 3 and Table 1 describe how to calculate the insertion loss budget if external cards are to be used. Figure 4 and Table 2 explain how to calculate the insertion loss budget if all the devices are on the carrier card.

For Serial Advanced Technology Attachment (SATA), there are three main considerations when designing carrier cards that use the SATA bus. First, all SATA signals must be referenced to ground, and layer changes must be minimized. To keep noise to a minimum, SATA traces should not be routed under crystals,

Segment	Loss (dB)	Notes
L_A	3.46	Allowance for 5.15 inches of module trace @ 0.28 dB / GHz / inch and 1.66 dB crosstalk allowance. Coupling caps not included.
Coupling Caps	1.19	From PCI Express Card Electromechanical Spec., Rev. 1.1, parameters (L_{ST} - L_{SR}). Includes crosstalk allowance of 0.79 dB.
L_B	0.25	COM Express connector at 1.25 GHz measured value.
L_C	4.40	Allowance for 9 inches of Carrier Board trace @ 0.28 dB / GHz / inch and a 1.25 dB crosstalk allowance.
L_D	1.25	PCI Express Card Electromechanical Spec Rev 1.1 "guard band" allowance for slot connector - includes 1.0 dB connector loss.
L_E	2.65	From PCI Express Card Electromechanical Spec., Rev. 1.1 (without coupling caps; L_{AR}). Implied crosstalk allowance is 1.25 dB.
Total	13.20	

Table 1 PCI Express Segment Losses (Figure 3).

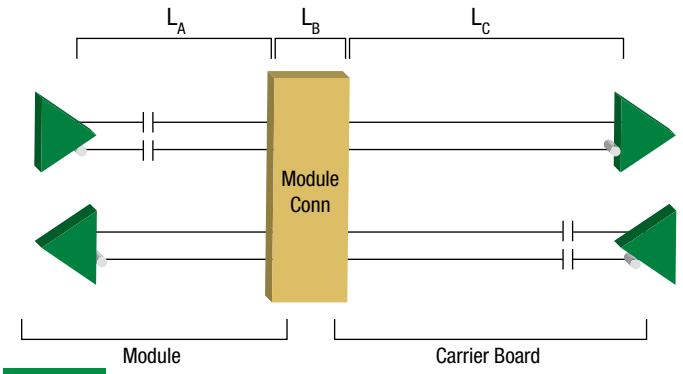

Figure 4 PCI Express insertion loss budget with carrier card PCIE device involves the module, module connector and carrier card equals the sum of $L_A+L_B+L_C$.

oscillators, clock drivers or other noise generators. Figure 5 and Table 3 explain the insertion loss budget for SATA applications.

Video and USB

USB is becoming ubiquitous and will definitely appear on COM carrier cards. There are also number of different video options for which the design considerations on the carrier cards can vary. Video options include SDVO, LVDS, TV Out and VGA.

In the case of serialized digital video output (SDVO) carrier cards, COM Express modules support two SDVO interfaces (SDVOB [channel 1] and SDVOC [channel 2]). Each SDVO interface is comprised of three differential pairs and associated differential clock. When designed for DVI mode, carrier cards must incorporate video CODEC for SDVO interfaces and DVI traces must be routed on the signal layer adjacent to a continuous ground plane.

AC coupling capacitors for each channel should be installed on the COM module avoiding the need for extra capacitors for AC coupling. When implementing SDVO devices on the carrier card, the SDVO_CLK and SDVO_DATA lines require pull-up resistors to 2.5V +/- 5% (resistor value should be 3.5K ohms). Therefore, when implementing slots for SDVO carrier cards on the carrier card, there is no need to use pull-up resistors because add-on SDVO cards already have them.

When creating LVDS carrier cards, route the signal traces as strip line and keep them ground-referenced. Also, control differential impedance of the pairs within 100 ohms ±15% and length match them to within 20 mils. Two pairs of clocks (LVDS_A_CK+/-, LVDS_B_CK+/-) must also be length matched within 20 mils and each channel of the LVDS bus must be length matched with corresponding clocks within 20 mils. Minimal space between neighboring traces and differential pairs must be kept over 20 mils. LVDS flat panel differential pairs (LVDS_A+/-, LVDS_B+/-. LVDS_A_CK+/-, LVDS_B_CK+/-) shall have 100 ohm terminations across the pairs at the destination. Unused LVDS lines may be left open.

If TV Out is used, the TV_DAC_A, TV_DAC_B and TV_DAC_C lines must be terminated on the carrier card through a 150-ohm resistor to ground. Termination resistors must be placed close to the external TV-Out connector(s). These lines may be left unterminated if the TV Out function is not used. If analog VGA is used, terminate the VGA_RED, VGA_GRN and VGA_BLU signals on the carrier card through a 150-ohm resistor to ground. Resistors should be placed close to the VGA connector on the carrier card. However, these lines may be left unterminated when the analog VGA function is not used.

Developing a USB carrier card becomes easier—for one thing, USB pairs require no termination. Still there are several

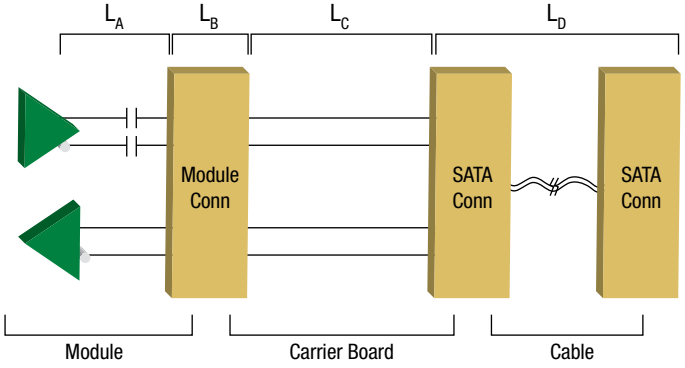

Figure 5 SATA insertion loss budget for module, carrier card, connectors and cable (sum of $L_A+L_B+L_C+L_D$).

Segment	Loss (dB)	Notes
L_A	3.46	Allowance for 5.15 inches of module trace @ 0.28 dB / GHz / inch and 1.66 dB crosstalk allowance. Coupling caps not included.
Coupling Caps	1.19	From PCI Express Card Electromechanical Spec., Rev. 1.1, parameters (L_{ST} - L_{SR}). Includes crosstalk allowance of 0.79 dB.
L_B	0.25	COM Express connector at 1.25 GHz measured value.
L_C	8.30	Allowance for 15.85 inches of Carrier Board trace @ 0.28 dB / GHz / inch and a 2.75 dB crosstalk allowance.
Total	13.20	

Table 2 PCI Express Insertion Loss Budget.

Segment	Loss (dB)	Notes
L_A	Gen 1: 1.26 Gen 2: 1.68	Up to 3.0 inches of module trace at 0.28 dB / GHz / inch Up to 2.0 inches of module trace at 0.28 dB / GHz / inch
Coupling Caps	Gen 1: 0.40 Gen 2: 0.40	
L_B	Gen 1: 0.25 Gen 2: 0.38	COM Express connector at 1.5 GHz measured value COM Express connector at 3.0 GHz measured value
L_C	Gen 1: 3.07 Gen 2: 2.52	Up to 7.2 inches of carrier card trace at 0.28 dB / GHz / inch Up to 3.0 inches of carrier card trace at 0.28 dB / GHz / inch
L_D	Gen 1: 6.0 Gen 2: 6.0	Source specification cable and cable connector allowance
Total	Gen 1: 10.98 Gen 2: 10.98	

Table 3 SATA Gen1 (1.5 GHz) and Gen 2 (3.0 GHz) Insertion Loss Budget Calculated.

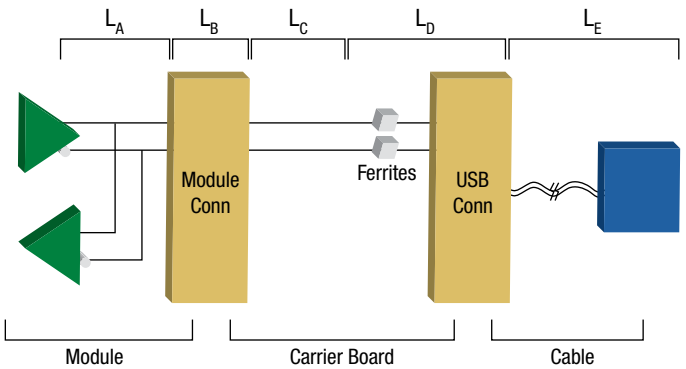

Figure 6 USB insertion loss budget includes module, carrier card and cable (sum of $L_A+L_B+L_C+L_D+L_E$).

Segment	Loss (dB)	Notes
L_A	0.67	Up to 6 inches of module trace @ 0.28 dB / GHz / inch
L_B	0.05	COM Express connector at 400 MHz measured value.
L_C	1.68	Up to 14 inches of Carrier Board trace @ 0.28 dB / GHz / inch
L_D	1.00	USB Connector and ferrite loss
L_E	5.80	USB cable and far end connector loss, per source specification
Total	9.20	

Table 4 USB Loss Budget Calculated.

considerations. USB supplies its own current and the signals USB_0_1_OC#, USB_2_3_OC#, USB_4_5_OC# and USB_6_7_OC# are used to flag a USB over-current situation. Carrier card USB current monitors may pull these lines to ground with open drain drivers to indicate that the monitor's current limit has been exceeded. Figure 6 and Table 4 explain how to calculate USB insertion loss budget for the module, carrier card, connectors and cable to stay within specified limits.

COM Express enables decoupling the processor module from the rest of the system, which allows manufacturers and developers to focus their attention on designing I/O carrier cards that can address the various applications, product configurations, price points and models specific to their industry. A multiplatform architecture using COM Express for embedded design allows the reuse and interchangeability of modules and/or carriers and permits use of common software libraries, operating system support packages and application middleware. This approach reduces design risk, as the most difficult piece of the design is already done, cuts development expenses and therefore provides faster time-to-market and time-to-revenue.

RadiSys
Hillsboro, OR.
(503) 615-1100.
[www.radisys.com].

RED ROCK TECHNOLOGIES
Mass Storage Modules
for VMEbus and CompactPCI®
NEW Product!
DVD-RW / CD-RW / CDROM
in VMEbus Form Factor
with Ultra Wide SCSI LVD interface.
See the full line of VMEbus single and
multi-slot mass storage module products at
www.RedRockTech.com
or call Toll-Free: 800-808-7837
Red Rock Technologies, Inc. — 480-483-3777

PC/104 SHOWCASE

Featuring the latest in PC/104 technology

CoreModule™ 800

PCI-104 Single Board Computer

- RoHS-Compliant
- Rugged, ultra-high performance compact module
- Low-power Centrino® performance
- -20° to + 70°C standard operating temperature
- PCI-104 without wings
- Advanced networking, high performance graphics

Ampro Computers, Inc.
Phone: (408) 360-0200
Fax: (408) 360-0222
E-mail: info@ampro.com
Web: www.ampro.com

EMERALD-MM-8P

- 8 serial ports, RS-232/RS-422/RS-485, 460.8kbps
- Programmable protocols, I/O addresses, and interrupt levels
- 16C654 UARTS with 64-byte FIFOs and interrupt sharing
- EEPROM storage of configuration data for automatic power-on configuration
- 8 programmable digital I/O lines
- -40°C to +85°C operation

Diamond Systems Corporation
Phone: (800) 36-PC104
Fax: (650) 810-2525
E-mail: sales@diamondsystems.com
Web: www.diamondsystems.com

AMPRO LittleBoard™ 800

EBX Single Board System

- World's first and only rugged EBX Pentium® M SBC
- Ultra-high performance, high speed I/O and full PC/104-Plus expandability
- -40° to +85°C fanless for reliability; conformal coating option
- 50% thicker PCB
- 3 year warranty – 5+ year life cycle policy
- OS support includes Windows® XP, XPe, CE, Linux, QNX and VX Works

Ampro Computers, Inc.
Phone: (408) 360-0200
Fax: (408) 360-0222
E-mail: info@ampro.com
Web: www.ampro.com

DIAMOND-MM-32X

- 32 analog inputs, 16-bit A/D resolution
- 250KHz A/D sampling rate
- Programmable input ranges
- 1024-sample A/D FIFO, programmable threshold
- Automatic autocalibration for care-free, high accuracy measurements
- 4 analog outputs, 12-bit resolution
- D/A waveform generator
- 24 digital I/O, programmable
- Counter/timers for A/D and D/A timing
- -40 to +85°C operation

Diamond Systems Corporation
Phone: (800) 36-PC104
Fax: (650) 810-2525
E-mail: sales@diamondsystems.com
Web: www.diamondsystems.com

Q104-1553

High Density PC/104-*Plus* Interface

- One, two, or four dual-redundant MIL-STD-1553 channels
- Multi-function (simultaneous BC, 31 RTs, and BM)
- Windows XP, 2000, VxWorks, Linux, and other OS available
- 1 MB RAM per channel
- PC/104-*Plus* or regular PC/104 formats

Condor Engineering
Phone: (805) 965-8000
Fax: (805) 963-9630
E-mail: sales@condoreng.com
Web: www.condoreng.com

SBC1670

XScale PC/104 Computer with LCD Panel Support

- 520MHz Low-power ARM processor
- 800 x 600 Color LCD interface
- 10/100BASE-T Ethernet
- Five serial ports
- Debounced keypad interface
- 128MB SDRAM, 64MB Flash
- CompactFlash connector
- PC/104 expansion
- -40°C to +85°C operation available
- Linux, Windows CE, VxWorks®

Micro/sys, Inc.
Phone: (818) 244-4600
Fax: (818) 244-4246
E-mail: info@embeddedsys.com
Web: www.embeddedsys.com

PC/104 SHOWCASE

SBC1586ET

Pentium PC/104 with CompactFlash, -40°C to +85°C operation

- Low-Power Pentium, 166 or 266MHz
- Up to 256MB SDRAM
- CompactFlash connector
- 4 COM ports, LPT1, KBD, mouse, USB, IDE, RTC, watchdog
- 10/100BASE-T Ethernet
- PC/104 expansion
- DOS, Linux, CE, NT, VxWorks®, RTOS

Micro/sys, Inc.
Phone: (818) 244-4600
Fax: (818) 244-4246
E-mail: info@embeddedsys.com
Web: www.embeddedsys.com

Lynx PC/104-Plus Single Board Computer

- AMD Élan SC520 x86 processor
- 64 MB Soldered-on RAM
- 4 COM ports (two RS-232 / two RS-422/485)
- 2 general purpose timers
- IDE interface, LPT port, PS/2 mouse / keyboard
- 10/100 Ethernet port
- Rugged design and quality manufacturing
- Long-Term Availability Guarantee

VersaLogic Corp.
Phone: (800) 824-3163
Fax: (541) 485-5712
E-mail: Info@VersaLogic.com
Web: www.VersaLogic.com

SBC1495ET

486/586 PC/104 wtih CompactFlash and Flat Panel

- Ready to run 486/586 computer, 120 or 133MHz
- 64MB SDRAM
- CompactFlash connector
- 2 COM ports, LPT1, KBD, mouse, USB, SVGA, IDE, RTC, watchdog
- 10/100BASE-T Ethernet
- PC/104 expansion
- -40° C to +85° C operation
- DOS, Linux, VxWorks®, WinCE/98, MSDOS

Micro/sys, Inc.
Phone: (818) 244-4600
Fax: (818) 244-4246
E-mail: info@embeddedsys.com
Web: www.embeddedsys.com

Jaguar PC/104-Plus Single Board Computer

- Intel Pentium III processor (Celeron Versions Available)
- High Performance graphics
- 2 COM ports (two RS-232 / two RS-422/485)
- IDE interface, LPT port, PS/2 mouse / keyboard
- 10/100 Ethernet port
- Rugged design and quality manufacturing
- Long-Term Availability Guarantee

VersaLogic Corp.
Phone: (800) 824-3163
Fax: (541) 485-5712
E-mail: Info@VersaLogic.com
Web: www.VersaLogic.com

MPC555

Analog, Digital, and Timer I/O for PC/104

- Sixteen 12-bit channels of analog input
- 100kHz A/D sampling rate
- Software selectable A/D input ranges
- 24 digital I/O lines
- 4 input jumpers and 8 LEDs
- Three 16-bit counter/timers
- Optional 12-bit DACs, 4 or 8 channels
- 5V-only operation
- Extended temp operation available

Micro/sys, Inc.
Phone: (818) 244-4600
Fax: (818) 244-4246
E-mail: info@embeddedsys.com
Web: www.embeddedsys.com

Gecko EPIC Single Board Computer

- EPIC form factor with PC/104-*Plus* expansion
- AMD GX500 processor
- Integrated graphics
- 4 USB ports, 4 COM ports (two RS-232 / two RS-422/485)
- IDE interface, PS/2 mouse and keyboard
- 10/100 Ethernet port
- Rugged design and quality manufacturing
- Long-Term Availability Guarantee

VersaLogic Corp.
Phone: (800) 824-3163
Fax: (541) 485-5712
E-mail: Info@VersaLogic.com
Web: www.VersaLogic.com

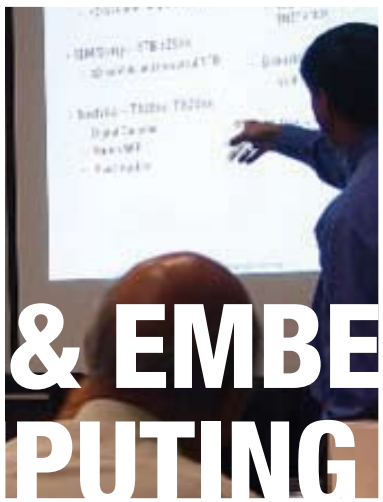

REAL-TIME & EMBEDDED COMPUTING CONFERENCE

Real-Time & Embedded computing conference

Enter a World of Embedded Computing Solutions

Register now at www.rtecc.com

Seminars

Through open-door technical sessions especially for those developing computer systems and time-critical applications. Get ahead with sessions on Embedded Linux, VME, PCI Express, ATCA, DSP, FPGA, Java, RTOS, SwitchFabric Interconnects, Windows, Wireless Connectivity, and much more.

Exhibits

Arranged in a unique setting to talk face-to-face with technical experts. Table-top exhibits make it easy to compare technologies, ask probing questions and discover insights that will make a big difference in your embedded computing world.

Exhibitor Spotlight

AAEON · Accelerated Technology · coverity · ENEA · Green Hills Software, Inc. · LYNUXWORKS™ · QNX · SBS Technologies · WDL Systems · WIND RIVER

2006 Locations

Melbourne	Huntsville	Milan	Rome	Albuquerque	Phoenix
February 28	**March 2**	**March 21**	**March 23**	**March 28**	**March 30**
Greenbelt	Houston	Calgary	Bristol	Beijing
Boston	Dallas	Vancouver BC	Patuxent River	Montreal
Minneapolis	Denver	San Diego	Taipei	Ottawa
Chicago	Salt Lake City	Los Angeles	Washington DC	Seattle
Helsinki	Detroit	Lucerne	Shenzhen	Portland
Oslo	Toronto	Eindhoven	Shanghai	

www.rtecc.com

Module-Based Design Platform Includes FPGAs

Small form-factor modules combined with FPGA-based IP offer a reconfigurable system-level design platform. Flexibility is increased by a sourcing community, or ecosystem, from which developers select components for upgrades, migration or new designs.

by Chris Wright
Ultimodule

The shrinking physical size of small form-factor modules offers new ways to support system-level product design. It is no longer efficient for designers to begin a design with an SBC or computer-on-module (COM) with an off-the-shelf operating system, and then find, evaluate and integrate a disparate collection of hardware and software add-ons.

A better solution combines small form-factor embedded modules with pre-integrated software and FPGA-based IP from various qualified technical experts. This combination can deliver a robust, reconfigurable system-level design platform. The solution supports fast project completion and provides an environment in which a developer selects system elements from a sourcing community, or ecosystem, for upgrades, migration or new designs.

Using an ecosystem of partners to support a common platform reduces the overall effort to source and qualify vendors and subsystem components. This saves organizational resources across engineering, purchasing, accounting and QA. By working together, ecosystem partners can create an enhanced common solution that they sell through combined channels, ultimately reaching more markets. They can also use each other to share in customer support and product maintenance activity.

There are two distinct advantages for designers who develop products with a pre-integrated embedded system design platform using FPGA-based small form-factor modules. The first advantage is the availability of pre-integrated software provided as optimized application engines that are designed to support the hardware module/CPU/OS combination. The second advantage is the availability of reconfigurable FPGA-based IP, which is delivered as pre-tested, application-specific reference designs tuned to work with the optimized software application engines.

When a design platform combines these two advantages, designers can efficiently tune a software engine to drive a specific FPGA-based IP block. By closely coupling the engines and support software to specific IP blocks, designers can easily swap in new IP blocks along with new application software to support upgrades or to change the product's functionality (Figure 1). This coupling of software to IP creates a flexible base set of building blocks that can be implemented as a single downloadable FPGA image. This system platform provides an efficient method for implementing an application-specific solution that can be further customized using pre-tested software engines.

Software: Beyond the OS

The software requirements of a robust embedded system platform must start with a basic OS, RTOS or otherwise. Whether the OS is a mature open-system product such as eCOS, or a commercial product such as Windows CE, basic network services such as FTP, PPP and HTTP Web server are required. In addition, features such as file systems and remote board firmware upgradeability are becoming standard requirements. Fast and reliable network queries, management and updates have become a given in the current checklist of platform requirements as the human resources to physically service or upgrade deployed products become more expensive and scarce.

Designers now require a lot more native capability in a system-level platform just to start their designs than were needed only a few years ago. There is a rising demand for readily available, optimized software engines or optional software add-ons in many of today's embedded systems. This software is essential for

Get Connected with companies mentioned in this article.
www.rtcmagazine.com/getconnected

fast, error-free project implementation. To establish a competitive advantage, a system-level platform provider must offer optional application-specific software libraries with optimized code tuned to execute on a selected module/CPU/OS combination. These engines reduce time-to-market by providing standard end-user APIs and protocols.

Application software and middleware are now more frequently pre-integrated into the final end product. Components like data acquisition software are being combined with programmable logic controller (PLC) functionality; Internet connectivity, including industry standard Web Services with well-defined XML application interfaces; and intelligent displays with built-in human-machine interface (HMI) programming capability. Providing these features requires the integration of many software components from different disciplines.

The greater benefit now resides in the pre-integration of software subsystems using the services of partners within an ecosystem that have the focus and expertise to bring design detail and field-testing to a solution. The cornerstone of the relationship among ecosystem partners includes cross-licensing and cross-selling of subsystem components, and the joint marketing and management of a shared hardware/software component repository.

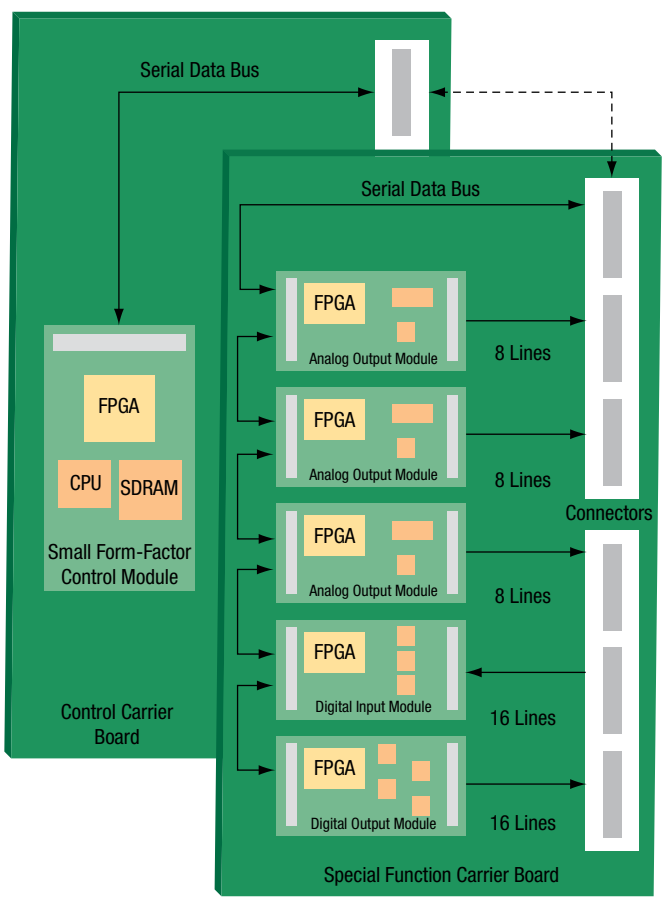

Figure 1 Through a serial data link, software engines in Ultimodule's Small Form-Factor Control Module on the Control Carrier Board target state machines in FPGA hardware on special function modules on the Special Function Carrier Board. Each module performs a separate, specialized function, allowing the designer to customize a system by swapping out modules.

Customizing the System Platform

A system-level platform optimized around a small form-factor module can be configured in many different ways. Two possible configurations are one that supports networking and data acquisition, and one that uses the same module reconfigured as an HMI platform that controls an operator display with touch control.

The networking and data acquisition platform (Figure 2) serves as a bridge between various communication protocols and I/O modules. To support this platform, several software libraries have been combined to create a multi-function software system that collects data and manages access to a network, using a local TCP/IP connection, or a Fieldbus and protocols like CAN or modBus. On this platform, the OS is not the differentiator. Instead, the differentiator is the pre-tested software that features a data acquisition library, a suite of available industrial protocols and Web services accessed though well-defined interface specifications.

The data acquisition library is at the center of this platform. It provides both data collection and control operations for applications in industrial and consumer appliances, or ad hoc controllers. Sitek, one of the partners in Ultimodule's ecosystem, developed this library and continuously supports it in the field. To serve the industrial automation markets in Europe and the United States, 150 industrial protocols are available with this platform.

Market segment-specific development tools are needed to program the finished product. In the industrial automation market, a suite of software control languages defined in the international standard IEC 61131-3 serve as the programming environment. Integration of the tool chain with downloadable core control software facilitates the rapid development of industrial control programs. Other partners in this ecosystem that deliver pre-integrated solutions are Smart Software Solutions (S3) with CoDeSys, and ICS Triplex with ISaGRAF.

The second configuration, the HMI control platform (Figure 3), is a system-level reference design that controls operator panels. It features configurable controllers and several pre-integrated software engines that support a wide range of LCD displays and touch control panels. The key software add-on is an object-oriented embedded scalable vector graphics (eSVG)-based visualization library and application software that provides graphical interfaces an operator uses to control systems, I/O communications and additional hardware in the HMI platform. Intesis is an ecosystem partner that developed this library and pre-integrates it in the reference design HMI platform.

The eSVG-HMI software engine satisfies most of the SVG Tiny and Basic profile features in the industry standard SVG 1.1 and SVG Mobile specifications. It also adheres to the Document Object Model (DOM) Level 2 Core Specification. DOM defines a set of core interfaces to create and manipulate the structure and contents of a document displayed on the LCD panel.

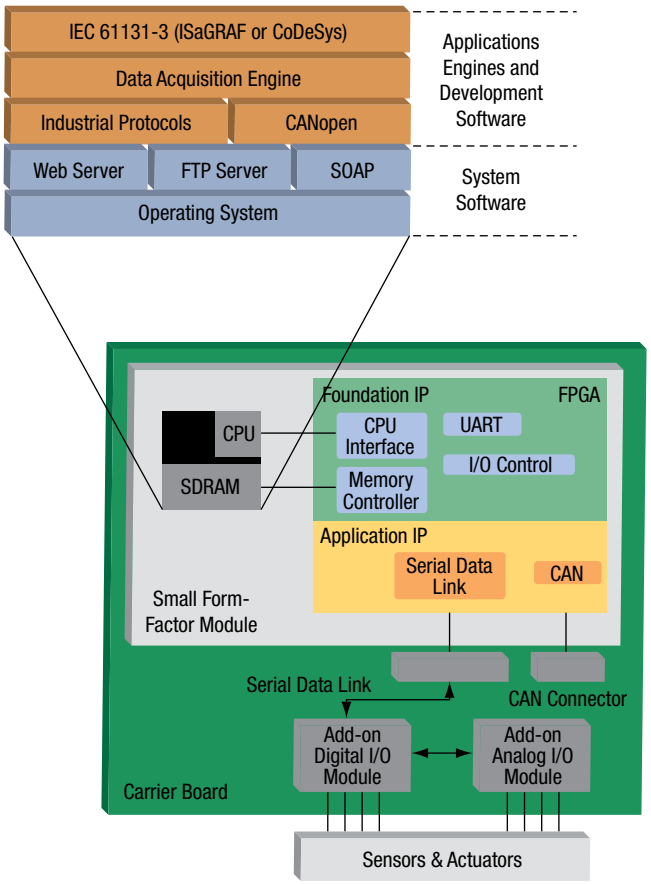

Figure 2 The software blocks required to create a Networking and Data Acquisition reference design using small form-factor modules. This reference design includes data acquisition software and CANopen, which are used to control the Serial Data Link and CAN application IP blocks.

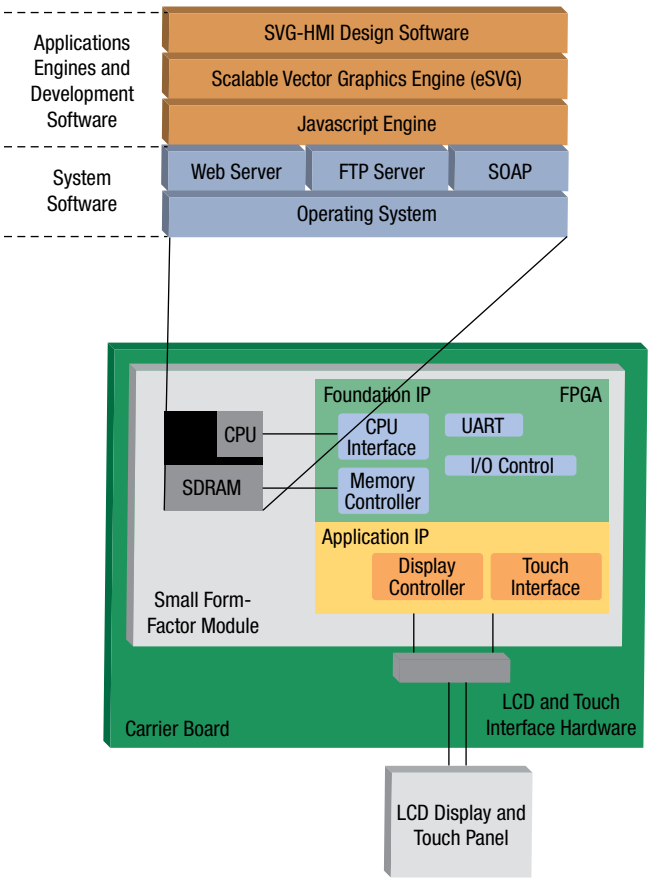

Figure 3 The software blocks required to create an HMI reference design using small form-factor modules. This reference design includes SVG software and engines that control the Display Controller and Touch Interface application IP block.

Creating an Application-Specific Reference Design

Using an FPGA as the nexus of a small form-factor module platform provides a means for compartmentalizing IP into a flexible solution, which offers greater versatility than a platform without an FPGA. One of the major benefits of using FPGA-based embedded modules is the ability to reuse pre-designed IP cores offered as IP blocks, which are optimized and packaged into a single downloadable image. When this image is downloaded into the FPGA, it changes the hardware available on the small form-factor module to meet the requirements of a specific product application. The IP blocks may be enhanced by including various software engines and APIs that add functionality to the platform.

Adding an FPGA to a small form-factor module means that selectable hardware peripheral components can be grouped in the FPGA as a set of IP core subsystems, thus creating a reference design. To use the platform, a designer simply selects the needed small form-factor board with a pre-installed mix of IP cores and adds some additional I/O modules to meet the application's requirements.

The essential elements of the platform are two types of IP block that are combined in the FPGA as a single bitstream. The first type is the foundation IP block and the second type is the application IP block. The foundation IP block includes a fundamental set of IP cores that provide the interfaces and peripherals to drive the module's basic operation. These cores include a memory controller, simple communications, timers and a CPU interface.

The application IP block includes special function cores that customize the FPGA to support a specific platform. It contains custom peripheral IP cores, such as video, graphics, matrix keyboard, touch controllers or a CAN 2.0 network interface. Combining the application IP block's cores with those of the foundation IP block creates an application-specific reference design.

Instead of creating new designs from scratch, a designer or integrator selects a reference design created from a repository of modular FPGA-based IP cores and relevant software engines. The reference design is pre-installed at the factory in the small form-factor module. To implement a product from the platform, a module with a preloaded reference design and its accompanying software is installed on a carrier board that provides the physical layer interface hardware needed to connect to the outside world.

The Power of Convergence...

Reliability. Scalability. Mobility...on a gateway blade.

Open architecture drives costs down and performance up.

Modular VoIP "blades" based around standard, interoperable modules like PMC and AdvancedMC reduce costs by limiting the number of unique blades that telecom OEMs and carriers have to purchase and stock. A softswitch or media gateway controller can be deployed in a minimal configuration and scaled up later (to OC-3 and beyond) without replacing the whole blade and without taking it offline. SBE provides high-performance DSP resource modules that deliver premium carrier class voice processing with world-class features using Texas Instruments' DSPs with Telogy Software. In addition, these modules support transcoding and transrating to enable the integration of voice, video, data, and wireless.

SBE products are scalable from daughterboard modules to complex gateway blades, and provide telecom carriers/service providers with a choice of programmable voice platforms featuring SBE's line of network interface cards, ranging from T1 and T3 to Gigabit Ethernet and IPsec/SSL/WLAN acceleration. Full Linux support is available on every board.

sbe

Linux On Demand

flexibility on demand | 925-355-2000 | info@sbei.com | www.sbei.com

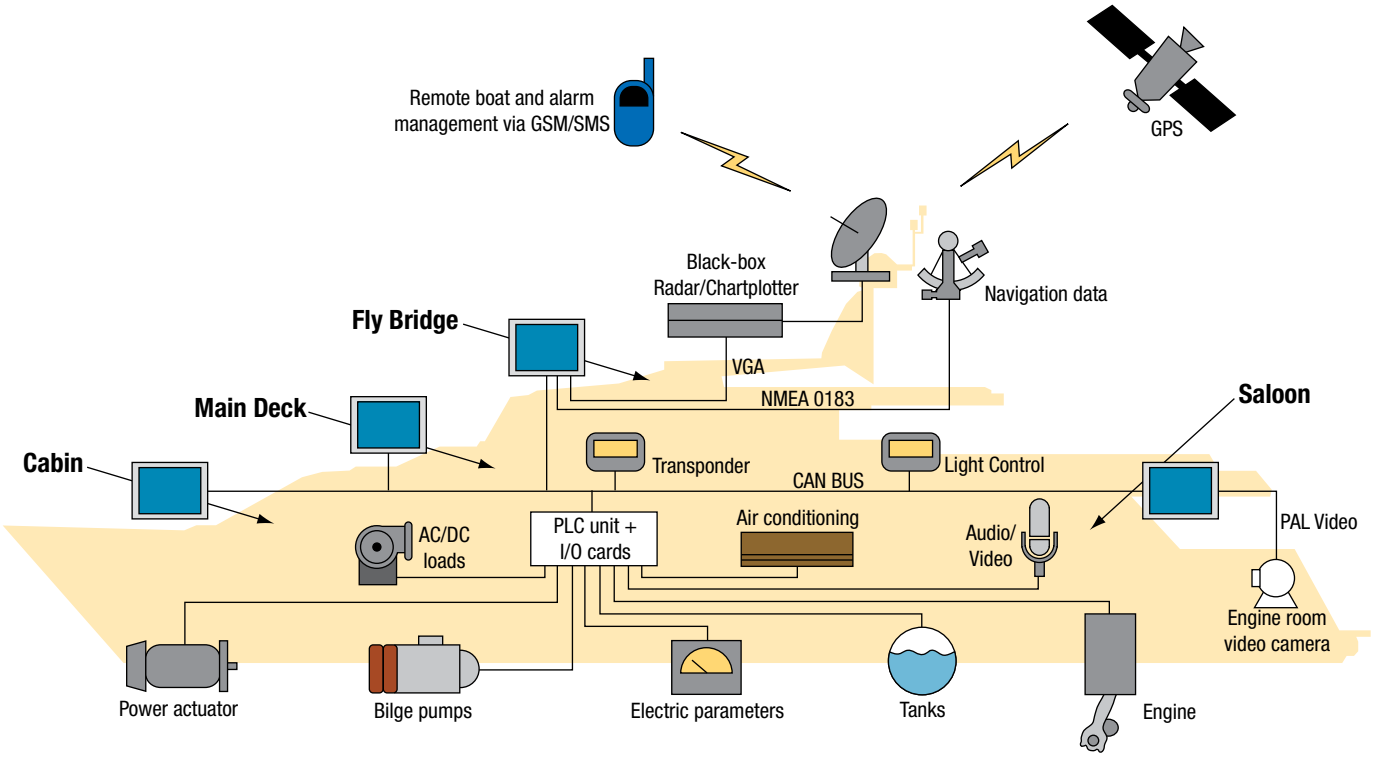

Figure 4 The integration of several HMI displays on the flying bridge, main deck, cabin and saloon of a luxury yacht. A networking and data acquisition platform based on special function I/O modules is used to control many of the yacht's onboard monitoring and control functions.

Combining Foundation IP and Application IP

The HMI platform reference design and the networking and data acquisition platform reference design both use the same foundation IP block, but different application IP blocks. The application IP block in the HMI reference design uses graphics and touch controller cores, while the application IP block in the networking and data acquisition reference design uses an FPGA-optimized serial data link core and a CAN controller core.

The key differentiating IP cores in the HMI control system are the graphics controller core and touch control panel interface. The graphics controller drives panel displays, such as electroluminescent and plasma, and LCD displays, such as STN, CSTN or TFT. The touch control panel is an interface to an A/D converter mapped through to the CPU. Intesis provides the eSVG- based HMI graphics application software that drives this display hardware. The eSVG-HMI visualization software engine executes code that controls the panel and displays graphic images of machine processes.

The key differentiating IP core in the networking and data acquisition system is the FPGA-optimized serial data link. This module-to-module data link uses a pair of communications IP cores to send command and data information between FPGA hardware subsystems in a multi-FPGA system. The data-link master core resides in the FPGA on the main small form-factor module, and a slave IP core resides in the FPGA on each remote I/O module.

This board-level data link extends the scope of the system by providing access to data collected by other IP cores spread across several FPGAs on I/O modules connected through the serial link. The data acquisition engine and I/O software drives the hardware and results in a system that can be easily configured for specific applications simply by plugging in the needed FPGA-based I/O modules.

Several products derived from these two reference designs are in production at NaviOP, a manufacturer of integrated systems for marine automation. The networking and data acquisition reference design is the basis of several systems that command and control all the onboard electric/electronic devices, including lights, pumps, winch, motor parameters and navigation instruments. NaviOP uses the HMI reference platform to enable onboard personnel to see the status and control the operation of these onboard devices (Figure 4).

Purchasing an SBC or COM with little more than a barebones OS, and then patching it together with a collection of hardware and software add-on solutions from independent vendors, is a high-risk development path. Instead, designers and integrators can use an FPGA-based, small form-factor module designed for advanced display control as the foundation for many different system-level platforms simply by selecting specific combinations of hardware IP and software components. A pre-integrated system-level platform based on open specifications, proven embedded hardware and software, and delivered in a cost-effective small form-factor, significantly raises the probability of completing the project on time and on budget. ◢

Ultimodule
Sunnyvale, CA.
(408) 734-2300.
[www.ultimodule.com].

Advertiser Index

Get Connected with technology and companies providing solutions now

Get Connected is a new resource for further exploration into products, technologies and companies. Whether your goal is to research the latest datasheet from a company, speak directly with an Application Engineer, or jump to a company's technical page, the goal of Get Connected is to put you in touch with the right resource. Whichever level of service you require for whatever type of technology, Get Connected will help you connect with the companies and products you are searching for.

www.rtcmagazine.com/getconnected

Company	Page	Website
ACCES I/O Products	50	www.accesio.com
Acromag	31,33,35	www.acromag.com
American Arium	16	www.arium.com
Ampro Computers, Inc.	24	www.ampro.com
Analog Devices, Inc.	18,19	www.analog.com/perform
ARCOM	58	www.arcom.com
BitMicro Networks, Inc.	6	www.bitmicro.com
BittWare	55	www.bittware.com
Digital Logic AG	60	www.adlogic-pc104.com
Diversified Technology	7	www.dtims.com
ELMA Electronic, Inc.	27	www.elma.com
Embedded Planet	36	www.embeddedplanet.com
Embedded Systems Conference	32	www.embedded.com/esc/sv
Hybricon Corporation	26	www.hybricon.com
Interactive Circuits and Systems	10	www.ics-ltd.com
Kontron America	8	www.kontron.com
Lippert Embedded Computers	43	www.lippert-at.com
Mercury Computer Systems	17	www.mercury.com
Micro Memory LLC	76	www.micromemory.com
Micro/sys, Inc.	51	www.embeddedsys.com
Microbus, Inc.	37	www.microbus-usa.com
Microsoft Tour	52	www.microsoftembedded.com
Microsoft Windows Embedded	38,39	
Muticore Expo	34	www.multicore-expo.com
Octagon Systems	2,3	www.octagonsystems.com
One Stop Systems	30	www.onestopsystems.com
PC/104 Showcase	66,67	
Performance Technologies	28	www.pt. com
Phoenix International	6	www.phenxint.com
QNX Software Systems, Ltd.	13	www.qnx.com
Real-Time & Embedded Computing Conference	68	www.rtecc.com
Red Rock Technologies, Inc.	65	www.redrocktech.com
RTC Digital Delivery	61	www.rtcmagazine.com
SanBlaze Technology	65	www.sanblaze.com
SBE, Inc.	72	www.sbei.com
SBS Technologies	4	www.sbs.com
Themis Computer	75	www.themis.com
Ultimate Solutions	47	www.ultsol.com
VadaTech	21	www.vadatech.com

RTC (Issn#1092-1524) magazine is published monthly at 905 Calle Amanecer, Ste. 250, San Clemente, CA 92673. Periodical postage paid at San Clemente and at additional mailing offices. POSTMASTER: Send address changes to RTC, 905 Calle Amanecer, Ste. 250, San Clemente, CA 92673.

Built to go in harm's way.

NEW! TPPC64™ (IBM PowerPC 970FX-based SBC)

6U VME64 engine with up to two IBM® 1.8 GHz PowerPC® 970FX CPUs
Up to 4GB of DDR400 SDRAM
VME interface - Universe II VME64x interface
Two 10/100/1000 Ethernet ports
Six USB ports
AC97 audio, two serial ports
Support for Linux™ OS
Ruggedized for up to 30G shock

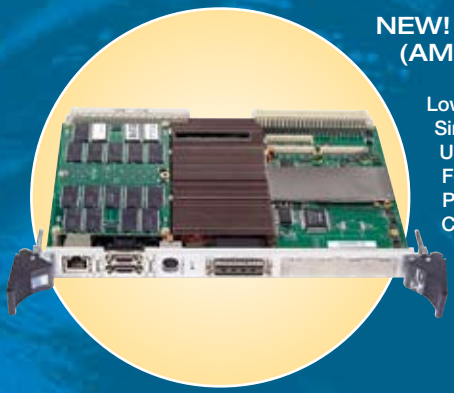

NEW! TA64™ (AMD Turion 64 Mobil-based SBC)

Low Power, AMD® Turion™ 64 Mobile CPU
Single/2-slot VME configurations
Up to 4GB ECC DDR memory
Four PMC slots, two USB ports
PS/2 keyboard and mouse port
Compatible with USPIIe-USB™ SBCs
64-bit Solaris 10, Windows® and Linux support
Ruggedized for up to 30G shock

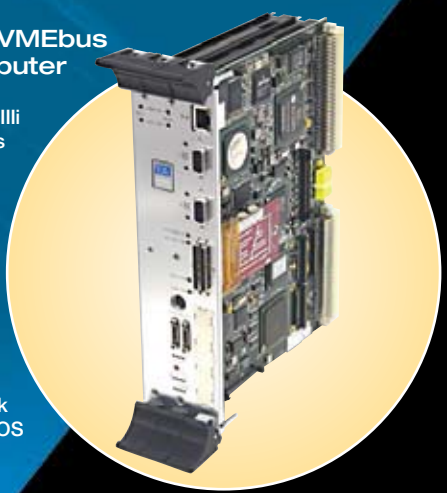

USP IIIi™ VMEbus Computer

Two UltraSPARC® IIIi 1.2 GHz CPUs
4GB 266DDR SDRAM per CPU, 8GB total
Optional TGA3D+™ graphics
Gigabit Ethernet port
One 64-bit/66-MHz PMC slot
Up to three PMC expansion slots
Dual FC-AL ports
Additional VME-backplane access
Up to 30G shock
Solaris™ OS

Our ultra-rugged, mission-critical SBCs now support more technologies than ever.

Extreme environments that need mission-critical reliability need Themis' single-board computers.

Our high performance, high availability computers have proven themselves in the harshest, most demanding applications and operating environments. Where other systems fail, our systems survive – in extreme temperatures and up to 30G shock.

And now Themis' family includes AMD and IBM PowerPC processors in addition to our well-known UltraSPARC VME and Compact PCI single board computers and graphic accelerators. So we can support applications in Solaris, Windows, Linux, UNIX®, and more.

All are designed for maximum flexibility and low Total Cost of Ownership. Because the budgeting process can be a pretty hostile environment, too.

So when your mission takes you in harm's way, you can rely on Themis.

www.themis.com (510) 252-0870

THEMIS

Transformational.

© 2005. Themis Computer, Themis, Themis logo, TA64, TGA3D+, USPIIe-USB and USPIIIi are trademarks or registered trademarks of Themis Computer. All other trademarks are property of their respective owners.

CoSine™ SoC

-on-

Othello®

VME CARRIERS FOR SWITCH FABRICS

MM-15x0 & MM-16x0

- Two CoSine 2VP70™ System-on-Chips
- Two XILINX® V-4 SX55 or LX160 CoSine Companion Devices
- Two XMC/PMC sites
- 14 independent memory arrays with total bandwidth of over 40 GB/s
- Four embedded PowerPC™ computers
- Complete on-board Serial RapidIO switch fabric connectivity
- VITA 41 or VPX VITA 46 formats
- Rugged Air Cooled and Conduction Cooled versions

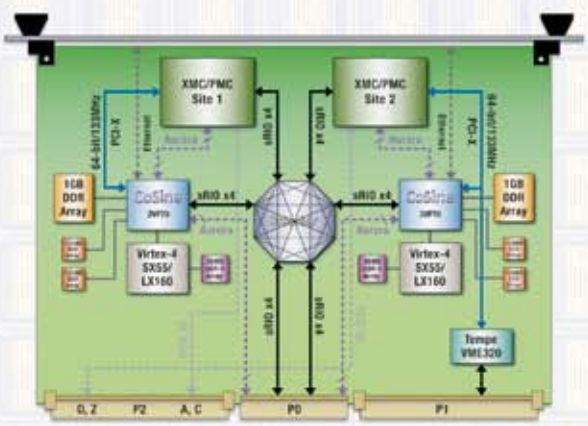

MM-15x0

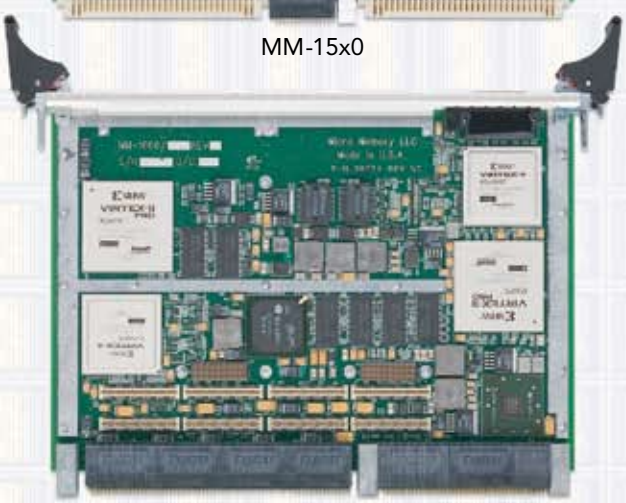

MM-16x0

Easily the most impressive board ever built for FPGA processing.

Quite possibly the most impressive ever built for VME.

RapidIO®

micro Memory

MICRO MEMORY
9540 Vassar Avenue, Chatsworth, California 91311 U.S.A.
(US) Tel 818 998 0070 • (US) Fax 818 998 4459
www.micromemory.com • sales@micromemory.com